Trace-Relating Compiler Correctness and Secure Compilation

Carmine Abate[1] Roberto Blanco[1] Ștefan Ciobâcă[2] Deepak Garg[3]
Cătălin Hrițcu[1] Marco Patrignani[4,5] Éric Tanter[6,1] Jérémy Thibault[1]

[1]Inria Paris [2]Al. I. Cuza University [3]MPI-SWS [4]Stanford [5]CISPA [6]University of Chile

Abstract. Compiler correctness is, in its simplest form, defined as the inclusion of the set of traces of the compiled program into the set of traces of the original program, which is equivalent to the preservation of all trace properties. Here traces collect, for instance, the externally observable events of each execution. This definition requires, however, the set of traces of the source and target languages to be exactly the same, which is not the case when the languages are far apart or when observations are fine grained. To overcome this issue, we study a generalized compiler correctness definition, which uses source and target traces drawn from potentially different sets and connected by an arbitrary relation. We set out to understand what guarantees this generalized compiler correctness definition gives us when instantiated with a non-trivial relation on traces. When this trace relation is not equality, it is no longer possible to preserve the trace properties of the source program unchanged. Instead, we provide a generic characterization of the target trace property ensured by correctly compiling a program that satisfies a given source property, and dually, of the source trace property one is required to show in order to obtain a certain target property for the compiled code. We show that this view on compiler correctness can naturally account for undefined behavior, resource exhaustion, different source and target values, side-channels, and various abstraction mismatches. Finally, we show that the same generalization also applies to many secure compilation definitions, which characterize the protection of a compiled program against linked adversarial code.

1 Introduction

Compiler correctness is an old idea [29, 31, 32] that has seen a significant revival in the recent past. This new wave was started by the creation of the CompCert verified C compiler [27] and continued by the proposal of many significant extensions and variants of CompCert [8, 17, 23, 24, 33, 39, 43, 44, 48] and the success of many other milestone compiler verification projects, including Vellvm [50], Pilsner [34], CakeML [45], Jasmin [4], CertiCoq [5], etc. Yet, even for these verified compilers, the precise statement of correctness matters. Since proof assistants are used to conduct the verification, an external observer does not have to understand the proofs in order to trust them, but one still has to deeply understand the statement that was proved. And this is true not just for correct compilation, but also for secure compilation, which is the more recent idea that our compilation chains should do more to also ensure security of our programs [3, 20].

Basic Compiler Correctness. The gold standard for compiler correctness is *semantic preservation*, which intuitively says that the semantics of a compiled program (in the target language) is compatible with the semantics of the original program (in the source language). For practical verified compilers, such as CompCert [27] and CakeML [45],

semantic preservation is stated extrinsically, by referring to *traces*. In these two settings, a trace is an ordered sequence of events—such as inputs from and outputs to an external environment—that are produced by the execution of a program.

A basic definition of compiler correctness can be given by the set inclusion of the traces of the compiled program into the traces of the original program, or formally [27]:

Definition 1.1 (Basic Compiler Correctness (CC)**).** *A compiler* $\downarrow$ *is* correct *iff*

$$\forall \mathsf{W}\ t.\ \mathsf{W}\!\downarrow\rightsquigarrow t \Rightarrow \mathsf{W}\rightsquigarrow t$$

This definition says that for any whole[1] source program W, if we compile it (denoted $\mathsf{W}\!\downarrow$), execute it with respect to the semantics of the target language, and observe a trace t, then the original W can produce *the same* trace t with respect to the semantics of the source language.[2] This definition is simple and easy to understand, since it only references a few familiar concepts: a compiler between a source and a target language, each equipped with a trace-producing semantics.

Beyond Basic Compiler Correctness. This basic compiler correctness definition assumes that any trace produced by a compiled program can be produced by the source program. This is a very strict requirement, and in particular implies that the source and target traces are drawn from the same set and that the same source trace corresponds to a given target trace. These assumptions are often too strong, and hence verified compiler efforts use different formulations of compiler correctness:

CompCert [27] The original compiler correctness theorem of CompCert [27] can be seen as an instance of basic compiler correctness, but it did not provide any guarantees for programs that can exhibit undefined behavior [40]. As allowed by the C standard, such unsafe programs were not even considered to be in the source language, so were not quantified over. This has important practical implications, since undefined behavior often leads to exploitable security vulnerabilities [10, 18, 19] and serious confusion even among experienced C and C++ developers [26, 40, 46, 47]. As such, since 2010, CompCert provides an additional top-level correctness theorem[3] that better accounts for the presence of unsafe programs by providing guarantees for them up to the point when they encounter undefined behavior [40]. This new theorem goes beyond the basic correctness definition above, as a target trace need only correspond to a source trace *up to the occurrence* of undefined behavior in the source trace.

CakeML [45] Compiler correctness for CakeML accounts for memory exhaustion in target executions. Crucially, memory exhaustion events cannot occur in source traces, only in target traces. Hence, dually to CompCert, compiler correctness only requires source and target traces to coincide up to the occurrence of a memory exhaustion event in the target trace.

Trace-Relating Compiler Correctness. Generalized formalizations of compiler correctness like the ones above can be naturally expressed as instances of a uniform definition, which we call *trace-relating compiler correctness*. This generalizes basic compiler

[1] For simplicity, for now we ignore separate compilation and linking, returning to it in §5.

[2] As a typesetting convention, we use a blue, sans-serif font for source elements, an **orange, bold** font for **target** ones and a *black, italic* font for elements common to both languages.

[3] Stated at the top of the CompCert file `driver/Complements.v` and discussed by Regehr [40].

correctness by (a) considering that source and target traces belong to *possibly distinct* sets $\mathsf{Trace_S}$ and $\mathbf{Trace_T}$, and (b) being parameterized by an arbitrary *trace relation* $\sim$.

Definition 1.2 (Trace-Relating Compiler Correctness ($\mathsf{CC}^{\sim}$)). *A compiler* $\downarrow$ *is* correct *with respect to a trace relation* $\sim \subseteq \mathsf{Trace_S} \times \mathbf{Trace_T}$ *iff*

$$\forall \mathsf{W}. \forall \mathbf{t}.\ \mathsf{W}{\downarrow} \rightsquigarrow \mathbf{t} \Rightarrow \exists \mathsf{s} \sim \mathbf{t}.\ \mathsf{W} \rightsquigarrow \mathsf{s}$$

This definition requires that, for any target trace $\mathbf{t}$ produced by the compiled program $\mathsf{W}{\downarrow}$, there exist a source trace s that can be produced by the original program W and is *related* to $\mathbf{t}$ according to $\sim$ (i.e. $\mathsf{s} \sim \mathbf{t}$). By choosing the trace relation appropriately, one can recover the different notions of compiler correctness presented above:

Basic CC Take $\mathsf{s} \sim \mathbf{t}$ to be $\mathsf{s} = \mathbf{t}$. Trivially, the basic CC of Definition 1.1 is $\mathsf{CC}^{=}$.

CompCert Undefined behavior is modeled in CompCert as a trace-terminating event $Goes_wrong$ that can occur in any of the languages (source, target, and all intermediate languages), so for a given (composition of) phase(s), we have $\mathsf{Trace_S} = \mathbf{Trace_T}$. But the relation between source and target traces with which to instantiate $\mathsf{CC}^{\sim}$ to obtain CompCert's current theorem is:

$$\mathsf{s} \sim \mathbf{t} \iff \mathsf{s} = \mathbf{t} \vee (\exists m \leq \mathbf{t}.\ \mathsf{s} = m{\cdot}Goes_wrong)$$

A compiler satisfying $\mathsf{CC}^{\sim}$ for this trace relation can turn a source trace ending in undefined behavior $m{\cdot}Goes_wrong$ (where "$\cdot$" is concatenation) either into the same trace in the target (first disjunct), or into a target trace that starts with the prefix m but then continues *arbitrarily* (second disjunct, "$\leq$" is the prefix relation).

CakeML Here, target traces are sequences of symbols from an alphabet $\mathbf{\Sigma_T}$ that has a specific trace-terminating event, $\mathbf{Resource_limit_hit}$, which is not available in the source alphabet Σ_S (i.e. $\mathbf{\Sigma_T} = \Sigma_S \cup \{\mathbf{Resource_limit_hit}\}$. Then, the compiler correctness theorem of CakeML can be obtained by instantiating $\mathsf{CC}^{\sim}$ with the following $\sim$ relation:

$$\mathsf{s} \sim \mathbf{t} \iff \mathsf{s} = \mathbf{t} \vee (\exists m.\ m \leq \mathsf{s}.\ \mathbf{t} = m{\cdot}\mathbf{Resource_limit_hit})$$

The resulting $\mathsf{CC}^{\sim}$ instance relates a target trace ending with $\mathbf{Resource_limit_hit}$ after executing m to a source trace that first produces m and then continues in a way given by the semantics of the source program.

Beyond undefined behavior and resource exhaustion, there are many other practical uses for $\mathsf{CC}^{\sim}$: in this paper we show that it also accounts for differences between source and target values, for a single source output being turned into a series of target outputs, and for side-channels.

On the flip side, the compiler correctness statement and its implications can be more difficult to understand for $\mathsf{CC}^{\sim}$ than for $\mathsf{CC}^{=}$. The full implications of choosing a particular $\sim$ relation can be subtle. In fact, using a bad relation can make the compiler correctness statement trivial or unexpected. For instance, it should be easy to see that if one uses the total relation, which relates all source traces to all target ones, the $\mathsf{CC}^{\sim}$ property holds for every compiler, yet it might take one a bit more effort to understand that the same is true even for the following relation:

$$\mathsf{s} \sim \mathbf{t} \iff \exists \mathsf{W}. \mathsf{W} \rightsquigarrow \mathsf{s} \wedge \mathsf{W}{\downarrow} \rightsquigarrow \mathbf{t}$$

Reasoning About Trace Properties. To understand more about a particular $\mathsf{CC}^{\sim}$ instance, we propose to also look at how it preserves *trace properties*—defined as sets of

allowed traces [25]—from the source to the target. For instance, it is well known that $\mathsf{CC}^{=}$ is equivalent to the preservation of all trace properties (where $W \models \pi$ reads "W satisfies π" and stands for $\forall t.\ W \rightsquigarrow t \Rightarrow t \in \pi$):

$$\mathsf{CC}^{=} \iff \forall \pi \in 2^{\text{Trace}}\ \forall \mathsf{W}.\ \mathsf{W} \models \pi \Rightarrow \mathsf{W}\!\downarrow \models \pi$$

However, to the best of our knowledge, similar results have not been formulated for trace relations beyond equality, when it is no longer possible to preserve the trace properties of the source program unchanged. For trace-relating compiler correctness, where source and target traces can be drawn from different sets, and related by an arbitrary trace relation, there are two crucial questions to ask:

1. For a source trace property π_{S} of a program—established for instance by formal verification—what is the strongest target property that any $\mathsf{CC}^{\sim}$ compiler is guaranteed to ensure for the produced target program?
2. For a target trace property $\pi_{\mathbf{T}}$, what is the weakest source property we need to show of the original source program to obtain $\pi_{\mathbf{T}}$ for the result of any $\mathsf{CC}^{\sim}$ compiler?

Far from being mere hypothetical questions, they can help the developer of a verified compiler to better understand the compiler correctness theorem they are proving, and we expect that any user of such a compiler will need to ask either one or the other if they are to make use of that theorem. In this work we provide a simple and definitive answer to these questions, for any instance of $\mathsf{CC}^{\sim}$. We observe that any trace relation $\sim$ induces two *property mappings* $\tilde{\tau}$ and $\tilde{\sigma}$, which are functions mapping source properties to target ones ($\tilde{\tau}$ standing for "to target") and target properties to source ones ($\tilde{\sigma}$ standing for "to source"):

$$\tilde{\tau}(\pi_{\mathsf{S}}) = \{\mathbf{t} \mid \exists \mathsf{s}.\, \mathsf{s} \sim \mathbf{t} \wedge \mathsf{s} \in \pi_{\mathsf{S}}\} \qquad \tilde{\sigma}(\pi_{\mathbf{T}}) = \{\mathsf{s} \mid \forall \mathbf{t}.\, \mathsf{s} \sim \mathbf{t} \Rightarrow \mathbf{t} \in \pi_{\mathbf{T}}\}$$

$\tilde{\tau}$ answers the first question above by mapping a given source property π_{S} to the target property that contains all target traces for which *there exists a related source trace* that satisfies π_{S}. Dually, $\tilde{\sigma}$ answers the second question by mapping a given target property $\pi_{\mathbf{T}}$ to the source property that contains all source traces for which *all related target traces* satisfy $\pi_{\mathbf{T}}$. Formally, we show that $\tilde{\tau}$ and $\tilde{\sigma}$ form a Galois connection and introduce two new correct compilation definitions in terms of *trace property preservation* (TP): $\mathsf{TP}^{\tilde{\tau}}$ quantifies over all source trace properties and uses $\tilde{\tau}$ to obtain the corresponding target properties. $\mathsf{TP}^{\tilde{\sigma}}$ quantifies over all target trace properties and uses $\tilde{\sigma}$ to obtain the corresponding source properties. We prove that these two definitions are equivalent to $\mathsf{CC}^{\sim}$, yielding a novel trinitarian view of compiler correctness (Figure 1).

$$\begin{array}{ccc} & \forall \mathsf{W}.\ \forall \mathbf{t}.\ \mathsf{W}\!\downarrow \rightsquigarrow \mathbf{t} \Rightarrow \exists \mathsf{s} \sim \mathbf{t}.\ \mathsf{W} \rightsquigarrow \mathsf{s} & \\ & \equiv & \\ & \mathsf{CC}^{\sim} & \\ \forall \pi_{\mathbf{T}}.\ \forall \mathsf{W}.\ \mathsf{W} \models \tilde{\sigma}(\pi_{\mathbf{T}}) \Rightarrow \mathsf{W}\!\downarrow \models \pi_{\mathbf{T}} \ \equiv\ \mathsf{TP}^{\tilde{\sigma}} & \Longleftrightarrow & \mathsf{TP}^{\tilde{\tau}} \ \equiv\ \forall \pi_{\mathsf{S}}.\ \forall \mathsf{W}.\ \mathsf{W} \models \pi_{\mathsf{S}} \Rightarrow \mathsf{W}\!\downarrow \models \tilde{\tau}(\pi_{\mathsf{S}}) \end{array}$$

Fig. 1: The equivalent compiler correctness definitions forming our trinitarian view.

Contributions

- We propose a new trinitarian view of compiler correctness that accounts for a non-trivial relation between source and target traces. While, as discussed above, specific instances of the $\mathsf{CC}^{\sim}$ definition have already been used in practice, we seem to be the first to propose assessing the meaningfulness of $\mathsf{CC}^{\sim}$ instances in terms of how properties are preserved between the source and the target, and in particular by looking at the property mappings $\tilde{\sigma}$ and $\tilde{\tau}$ induced by the trace relation $\sim$. We prove that $\mathsf{CC}^{\sim}$, $\mathsf{TP}^{\tilde{\sigma}}$, and $\mathsf{TP}^{\tilde{\tau}}$ are equivalent for any trace relation, as illustrated in Figure 1 (§2.3). In the opposite direction, we show that any property mappings that form a Galois connection induce a trace relation so that an analogous equivalence holds (§2.4). Finally, we extend these results from the preservation of trace properties to the larger class of subset-closed hyperproperties (§2.5) (e.g., noninterference).
- We use $\mathsf{CC}^{\sim}$ compilers of various complexities to illustrate that our view on compiler correctness naturally accounts for undefined behavior (§3.1), resource exhaustion (§3.2), different source and target values (§3.3), and differences in the granularity of data and observable events (§3.4). We expect these ideas to apply to any other discrepancies between source and target traces. For each compiler we show how to choose the relation between source and target traces and how the induced property mappings preserve interesting trace properties and subset-closed hyperproperties. We look at the way particular $\tilde{\sigma}$ and $\tilde{\tau}$ work on different kinds of properties and how the produced property can be expressed for different kinds of traces.
- We analyze the impact of correct compilation on noninterference [16], showing what can still be preserved (and thus also what is lost) when target observations are finer than source ones, e.g., side-channel observations (§4). We formalize the guarantee obtained by correct compilation of a noninterfering program as an *abstract noninterference* [15], a weakening of target noninterference. Dually we identify a class of declassifications of target noninterference for which source reasoning is possible.
- Finally, we show that the trinitarian view also extends to a large class of *secure compilation* definitions, formally characterizing the protection of the compiled program against linked adversarial code (§5). For each secure compilation definition we again propose both a property-free characterization in the style of $\mathsf{CC}^{\sim}$, and two characterizations in terms of preserving a class of source or target properties satisfied against arbitrary adversarial contexts. The additional quantification over contexts allows for finer distinctions when considering different property classes, so we study mapping classes not only of trace properties and hyperproperties, but also of relational hyperproperties [2]. An example secure compiler accounting for a target that can produce additional observations that are not possible in the source illustrates this approach.

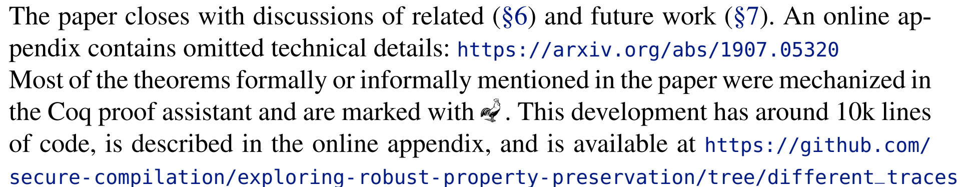

The paper closes with discussions of related (§6) and future work (§7). An online appendix contains omitted technical details: https://arxiv.org/abs/1907.05320
Most of the theorems formally or informally mentioned in the paper were mechanized in the Coq proof assistant and are marked with . This development has around 10k lines of code, is described in the online appendix, and is available at https://github.com/secure-compilation/exploring-robust-property-preservation/tree/different_traces

2 Trace-Relating Compiler Correctness

In this section, we start by generalizing the trace property preservation definitions from the end of the introduction to TP^{σ} and TP^{τ}, which depend on two *arbitrary* mappings σ and τ (§2.1). We prove that, whenever σ and τ form a Galois connection, TP^{σ} and TP^{τ} are equivalent (§2.2). We use this general result to close the trinitarian equivalence of Figure 1 for a given trace relation $\sim$ and its induced Galois connection $\tilde{\sigma} \leftrightarrows \tilde{\tau}$ (§2.3). This helps us assess the meaningfulness of a given trace relation by looking at the property mappings it induces. We also prove a dual result: for any σ and τ forming a Galois connection we can define a trace relation so that the trinitarian equivalence holds (§2.4). This allows us to construct new compiler correctness definitions starting from a desired mapping of properties. Finally, we generalize the classic result that compiler correctness (e.g., $\mathsf{CC}^{=}$) is enough to preserve not just trace properties but also all subset-closed hyperproperties [11]. For this we show that $\mathsf{CC}^{\sim}$ is also equivalent to subset-closed hyperproperty preservation, for which we also define both a version in terms of $\tilde{\sigma}$ and a version in terms of $\tilde{\tau}$ (§2.5).

2.1 Property Mappings

As explained in §1, trace-relating compiler correctness $\mathsf{CC}^{\sim}$, by itself, lacks a crisp description of which trace properties are preserved by compilation. Since even the syntax of traces can differ between source and target, one can either look at trace properties of the source (but then one needs to interpret them in the target), or at trace properties of the target (but then one needs to interpret them in the source).

Formally, these two interpretations can be seen as two property mappings:

$$\tau : 2^{\mathsf{Trace_S}} \to 2^{\mathbf{Trace_T}} \qquad \sigma : 2^{\mathbf{Trace_T}} \to 2^{\mathsf{Trace_S}}$$

For an arbitrary source program W, τ interprets a source property π_{S} as the *target guarantee* for $\mathsf{W}{\downarrow}$. Dually, σ defines a *source obligation* sufficient for the satisfaction of a target property $\boldsymbol{\pi_{\mathbf{T}}}$ after compilation. To be coherent with this informal interpretation, extra conditions on τ and σ seem natural (and as we will see in §2.2):

- First, the source property obtained by applying σ to the guarantee obtained by interpreting π_{S} in the target (via τ) should be weaker than π_{S} itself, i.e. $\sigma(\tau(\pi_{\mathsf{S}})) \supseteq \pi_{\mathsf{S}}$.
- Dually, the target property obtained via τ from the source obligation (via σ) for a target property $\boldsymbol{\pi_{\mathbf{T}}}$ should actually ensure that $\boldsymbol{\pi_{\mathbf{T}}}$ holds, i.e. $\tau(\sigma(\boldsymbol{\pi_{\mathbf{T}}})) \subseteq \boldsymbol{\pi_{\mathbf{T}}}$.

These two conditions on τ and σ are satisfied when the two maps form a *Galois connection* between the posets of source and target properties ordered by inclusion. Let us first recall the definition and the characteristic property of Galois connections [13, 30].

Definition 2.1 (Galois connection). *Let $(X, \preceq)$ and $(Y, \sqsubseteq)$ be two posets. A pair of maps, $\alpha : X \to Y$, $\gamma : Y \to X$ is a Galois connection iff it satisfies the following* adjunction law*: $\forall x \in X.\, \forall y \in Y.\, \alpha(x) \sqsubseteq y \iff x \preceq \gamma(y)$. α is referred to as the* lower adjoint *and γ as the* upper adjoint*. We will often write $\alpha : (X, \preceq) \leftrightarrows (Y, \sqsubseteq) : \gamma$ to denote a Galois connection, and simply $\alpha : X \leftrightarrows Y : \gamma$ when X and Y are powersets with inclusion, and even $\alpha \leftrightarrows \gamma$ when the involved sets are clear from context.*

Lemma 2.2 (Characteristic property of Galois connections). *If* α:$(X, \preceq) \leftrightarrows (Y, \sqsubseteq)$:$\gamma$ *is a Galois connection, then* α, γ *are monotone and they satisfy these properties:*

$$i) \quad \forall x \in X.\, x \preceq \gamma(\alpha(x)) \qquad\qquad ii) \quad \forall y \in Y.\, \alpha(\gamma(y)) \sqsubseteq y$$

If X, Y *are complete lattices, then* α *is continuous, i.e.* $\forall F \subseteq X.\, \alpha(\bigsqcup F) = \bigsqcup \alpha(F)$

These properties coincide with the conditions on property mappings introduced above, taking X to be source properties, Y to be target properties, α to be τ, γ to be and σ, and the order to be set inclusion, i.e. $\sigma(\tau(\pi_S)) \supseteq \pi_S$ and $\tau(\sigma(\pi_T)) \subseteq \pi_T$.

Perhaps surprisingly, the property mapping to the target (τ) corresponds to the abstraction function α, while the property mapping to the source (σ) corresponds to the concretization function γ. The explanation is that the definition of trace properties 2^{Trace} is contravariant in Trace, so while traces get more concrete when moving to the target, trace properties get, well, more abstract.

2.2 Revisiting Trace Property Preservation

We can now generalize trace property preservation (TP), relying on an interpretation of source properties in the target and target properties in the source.

Definition 2.3 (TP$^\sigma$ **and** TP$^\tau$). *Given two trace property mappings,* $\tau : 2^{\mathsf{Trace}_S} \to 2^{\mathbf{Trace}_T}$ *and* $\sigma : 2^{\mathbf{Trace}_T} \to 2^{\mathsf{Trace}_S}$*, for a compilation chain* $\cdot\!\downarrow$ *we define:*

$$\mathsf{TP}^\tau \equiv \forall \pi_S.\, \forall \mathsf{W}.\, \mathsf{W} \models \pi_S \Rightarrow \mathsf{W}\!\downarrow \models \tau(\pi_S) \qquad \mathsf{TP}^\sigma \equiv \forall \pi_T.\, \forall \mathsf{W}.\, \mathsf{W} \models \sigma(\pi_T) \Rightarrow \mathsf{W}\!\downarrow \models \pi_T$$

By uniqueness of adjoints [30], given a τ, there exists at most one σ such that the two form a Galois connection (and vice versa). Here, this means that if one is designing a compiler whose correctness goal is to ensure preservation of source properties according to a given mapping τ, then there can be only one corresponding property mapping σ from target properties to source properties (and vice versa). This suggests that, for any two property mappings defining a Galois connection, TP$^\tau$ and TP$^\sigma$ are equivalent.

Theorem 2.4 (TP$^\tau$ **and** TP$^\sigma$ **coincide** 🐓). *Let* $\tau : 2^{\mathsf{Trace}_S} \rightleftarrows 2^{\mathbf{Trace}_T} : \sigma$ *be a Galois connection, with* τ *and* σ *the lower and upper adjoints (resp.). Then* TP$^\tau$ $\iff$ TP$^\sigma$.

2.3 From Trace Relations to Property Mappings

We now investigate the relation between CC$^\sim$ and TP$^\tau$ and TP$^\sigma$. We show that, starting from a given trace relation $\sim$, it is always possible to define two property mappings $\tilde{\tau}$, $\tilde{\sigma}$ so that the three criteria are all equivalent. Dually, §2.4 will show that, given a Galois connection $\tau \rightleftarrows \sigma$, it is possible to define a relation that ensures the equivalence.

Definition 2.5 (Induced property mappings). *For an arbitrary trace relation* $\sim \subseteq \mathsf{Trace}_S \times \mathbf{Trace}_T$*, we define its* induced property mappings $\tilde{\tau}$ *and* $\tilde{\sigma}$ *as in §1:*

$$\tilde{\tau} = \lambda\, \pi_S.\ \{\mathsf{t} \mid \exists \mathsf{s}.\, \mathsf{s} \sim \mathsf{t} \wedge \mathsf{s} \in \pi_S\} \qquad \tilde{\sigma} = \lambda\, \pi_T.\ \{\mathsf{s} \mid \forall \mathsf{t}.\, \mathsf{s} \sim \mathsf{t} \Rightarrow \mathsf{t} \in \pi_T\}$$

These induced property mappings indeed form a Galois connection.

Theorem 2.6 (Induced adjunction). *$\tilde{\tau} \leftrightarrows \tilde{\sigma}$ is a Galois connection between the sets of trace properties of source and target traces ordered by set inclusion.*

We denote by $\mathsf{TP}^{\tilde{\tau}}$, $\mathsf{TP}^{\tilde{\sigma}}$ the criteria in Definition 2.3 for $\tilde{\tau}$ and $\tilde{\sigma}$ respectively. Theorem 2.7 below states the equivalence of $\mathsf{TP}^{\tilde{\tau}}$, $\mathsf{CC}^{\sim}$, and $\mathsf{TP}^{\tilde{\sigma}}$:

Theorem 2.7 ($\mathsf{TP}^{\tilde{\tau}}$, $\mathsf{CC}^{\sim}$, $\mathsf{TP}^{\tilde{\sigma}}$ coincide). *For any trace relation $\sim$ and induced property mappings $\tilde{\tau}$ and $\tilde{\sigma}$ (Definition 2.5), we have:* $\mathsf{TP}^{\tilde{\tau}} \iff \mathsf{CC}^{\sim} \iff \mathsf{TP}^{\tilde{\sigma}}$.

It follows that, for a $\mathsf{CC}^{\sim}$ compiler, $\tilde{\sigma}$ provides the source proof obligation *ensuring* a certain target property holds, and $\tilde{\tau}$ provides the target *guarantee* of a compiled program every time a property holds for the source program. Sometimes the lifted properties may be trivial: the target guarantee can be true, and the source obligation can be false.

2.4 From Property Mappings to Trace Relations

A compiler designer might start by first determining the target guarantees (τ), and then aim at showing TP^{τ}, which can, however, be challenging to prove directly. This section shows that, given a Galois connection between source and target properties, we can define a trace relation $\sim$ that ensures the equivalence between $\mathsf{CC}^{\sim}$, TP^{τ}, and TP^{σ}. Consequently, instead of proving TP^{τ} directly, it is sufficient to prove $\mathsf{CC}^{\sim}$, for which convenient proof techniques exist in the literature [7, 45].

Definition 2.8 (Induced trace relation ($\hat{\sim}$)). *Given a Galois connection $\tau : 2^{\mathsf{Trace_S}} \leftrightarrows 2^{\mathbf{Trace_T}} : \sigma$, we define the induced trace relation $\hat{\sim} \subseteq \mathsf{Trace_S} \times \mathbf{Trace_T}$ so that the target guarantee for π_S, $\tau(\pi_\mathsf{S})$, coincides with the image of the relation on the same π_S.*

$$\mathsf{s} \mathbin{\hat{\sim}} \mathbf{t} \iff \mathbf{t} \in \tau(\{\mathsf{s}\})$$

Notice that σ does not appear in the definition, but it is uniquely defined by being the upper adjoint of τ [30], so $\hat{\sim}$ depends on both τ and σ.

Theorem 2.9 (TP^{τ}, $\mathsf{CC}^{\hat{\sim}}$, and TP^{σ} all coincide). *Given a Galois connection $\tau \rightleftarrows \sigma$ between property mappings, and their induced trace relation $\hat{\sim}$, we have:*

$$\mathsf{TP}^{\tau} \iff \mathsf{CC}^{\hat{\sim}} \iff \mathsf{TP}^{\sigma}$$

2.5 Preservation of Subset-Closed Hyperproperties

A $\mathsf{CC}^{=}$ compiler ensures the preservation of not only trace properties but also all subset-closed hyperproperties, as this class of hyperproperties is known to be preserved by refinement [11]. An example of a subset-closed hyperproperty is *noninterference* [11], and a $\mathsf{CC}^{=}$ compiler guarantees that if W is noninterfering with respect to the inputs and outputs in the trace then so is $\mathsf{W}\downarrow$. To be able to talk about how properties such as noninterference are mapped here we propose another trinitarian view involving $\mathsf{CC}^{\sim}$ and preservation of subset-closed hyperproperties (Theorem 2.12), slightly weakened in that source and target property mappings will need to be closed under subsets.

First recall that a program satisfies a hyperproperty when its complete set of traces, which from now on we will call its *behavior*, is a member of the hyperproperty [11].

Definition 2.10 (Hyperproperty Satisfaction). *A program W satisfies a hyperproperty H, written $W \models H$, if and only if $beh(W) \in H$, where $beh(W) = \{t \mid W \rightsquigarrow t\}$.*

Hyperproperty preservation is a strong requirement in general. Fortunately many interesting hyperproperties are *subset-closed* (SCH for short), which simplifies preservation of such hyperproperties, since it suffices to show that the behaviors of the compiled program refine the behaviors of the source one, which coincides with the statement of $\mathsf{CC}^{=}$.

To talk about hyperproperty preservation in the trace-relating setting, we need an interpretation of source hyperproperties into the target and vice versa. The one we consider builds on top of the two trace property mappings τ and σ, which are naturally lifted to hyperproperty mappings. This way we are able to extract two hyperproperty mappings from a trace relation similarly to §2.3.

Definition 2.11 (Lifting property mappings to hyperproperty mappings). *Let $\tau : 2^{\mathsf{Trace_S}} \to 2^{\mathbf{Trace_T}}$ and $\sigma : 2^{\mathbf{Trace_T}} \to 2^{\mathsf{Trace_S}}$ be arbitrary property mappings. The images of $\mathsf{H_S} \in 2^{2^{\mathsf{Trace_S}}}$, $\mathbf{H_T} \in 2^{2^{\mathbf{Trace_T}}}$ under τ and σ respectively are:*

$$\tau(\mathsf{H_S}) = \{\tau(\pi_{\mathsf{S}}) \mid \pi_{\mathsf{S}} \in \mathsf{H_S}\} \qquad \sigma(\mathbf{H_T}) = \{\sigma(\boldsymbol{\pi}_{\mathbf{T}}) \mid \boldsymbol{\pi}_{\mathbf{T}} \in \mathbf{H_T}\}$$

Formally we are defining two new mappings, this time on hyperproperties, but by a small abuse of notation we still denote them by τ and σ.

Interestingly, it is not possible to apply the same argument used for $\mathsf{CC}^{=}$ to show that a $\mathsf{CC}^{\sim}$ compilation chain guarantees $\mathsf{W}{\downarrow} \models \tilde{\tau}(\mathsf{H_S})$ whenever $\mathsf{W} \models \mathsf{H_S}$. That proof breaks because the target hyperproperty $\tilde{\tau}(\mathsf{H_S})$ is not necessarily subset-closed, even if the original hyperproperty $\mathsf{H_S}$ is. In the following theorem we consider the loss of precision due to the two interpretations of hyperproperties, and close $\tilde{\tau}(\mathsf{H_S})$ under subsets (indicated as $Cl_{\subseteq}$); a similar treatment is given to $\tilde{\sigma}$ as well.

Theorem 2.12 (Preservation of Subset-Closed Hyperproperties 🐓). *For any trace relation $\sim$ and induced hyperproperty mappings $\tilde{\tau}$ and $\tilde{\sigma}$, for $Cl_{\subseteq}(H) = \{\pi \mid \exists \pi' \in H.\, \pi \subseteq \pi'\}$, we have: $\mathsf{SCHP}^{Cl_{\subseteq} \circ \tilde{\tau}} \iff \mathsf{CC}^{\sim} \iff \mathsf{SCHP}^{Cl_{\subseteq} \circ \tilde{\sigma}}$.*

$$\mathsf{SCHP}^{Cl_{\subseteq} \circ \tilde{\tau}} \equiv \forall \mathsf{W} \forall \mathsf{H_S} \in \mathsf{SCH_S}.\mathsf{W} \models \mathsf{H_S} \Rightarrow \mathsf{W}{\downarrow} \models Cl_{\subseteq}(\tilde{\tau}(\mathsf{H_S}))$$

$$\mathsf{SCHP}^{Cl_{\subseteq} \circ \tilde{\sigma}} \equiv \forall \mathsf{W} \forall \mathbf{H_T} \in \mathbf{SCH_T}.\mathsf{W} \models Cl_{\subseteq}(\tilde{\sigma}(\mathbf{H_T})) \Rightarrow \mathsf{W}{\downarrow} \models \mathbf{H_T}$$

In conclusion, $\mathsf{CC}^{\sim}$ can also be used to reason about the preservation of subset-closed hyperproperties, but one has to be aware of the potential loss of precision. We illustrate this in §4 when considering the preservation of noninterference in a $\mathsf{CC}^{\sim}$ compiler to a target language with additional observations.

3 Instances of Trace-Relating Compiler Correctness

The trace-relating view of compiler correctness of §2 can serve as a unifying framework for studying a range of interesting compilers. This section provides several representative instantiations of the framework: source languages with undefined behavior that compilation can turn into arbitrary target behavior (§3.1), target languages with resource exhaustion that cannot happen in the source (§3.2), changes in the representation of values (§3.3), and differences in the granularity of data and observable events (§3.4).

3.1 Undefined Behavior

In this section, we expand upon the discussion of undefined behavior from §1. We first stick to CompCert's model, where source and target alphabets are the same, including the event for undefined behavior. The relation allows relaxing equality by potentially replacing undefined behavior with an arbitrary sequence of events.

Example 3.1 (CompCert-like Undefined Behavior Relation). Source and target traces are sequences of events drawn from Σ, where $Goes_wrong \in \Sigma$. The symbol $Goes_wrong$ is a terminal event that represents an undefined behavior. We recall the trace relation for undefined behavior from the introduction:

$$\mathsf{s} \sim \mathbf{t} \iff \mathsf{s} = \mathbf{t} \vee \exists m \leq \mathbf{t}.\ \mathsf{s} = m \cdot Goes_wrong$$

Each trace of a target program produced by a $\mathsf{CC}^{\sim}$ compiler is either also a trace of the original source program, or it has finite prefix that the source program also produces, immediately before encountering undefined behavior. One of CompCert's correctness theorem can be rephrased as this variant of $\mathsf{CC}^{\sim}$.

Trace properties We proved that the two property mappings induced by the relation can be written as follows (🐓):

$$\begin{aligned}\tilde{\sigma}(\boldsymbol{\pi_{\mathbf{T}}}) &= \{\mathsf{s} \mid \mathsf{s} \in \boldsymbol{\pi_{\mathbf{T}}} \wedge \mathsf{s} \neq m \cdot Goes_wrong\} \\ &\quad \cup \{m \cdot Goes_wrong \mid \forall \mathbf{t}, m \leq \mathbf{t} \implies \mathbf{t} \in \boldsymbol{\pi_{\mathbf{T}}}\} \\ \tilde{\tau}(\pi_{\mathsf{S}}) &= \{\mathbf{t} \mid \mathbf{t} \in \pi_{\mathsf{S}}\} \cup \{\mathbf{t} \mid \exists m \leq \mathbf{t}.\ m \cdot Goes_wrong \in \pi_{\mathsf{S}}\}\end{aligned}$$

These two mappings explain what a $\mathsf{CC}^{\sim}$ compiler ensures for the $\sim$ relation above:

– The target-to-source mapping $\tilde{\sigma}$ states that to prove that a compiled program has a property $\boldsymbol{\pi_{T}}$ using source-level reasoning, one has to prove that any trace produced by the source program must either be a target trace satisfying $\boldsymbol{\pi_{T}}$ or have undefined behavior provided that *any continuation* of the trace substituted for the undefined behavior satisfies $\boldsymbol{\pi_{T}}$.

– The source-to-target mapping $\tilde{\tau}$ states that by compiling a program satisfying a property π_{S} we obtain a program that produces traces that satisfy the same property or that extend a source trace that ends in undefined behavior.

These definitions can help us reason about programs. For instance, $\tilde{\sigma}$ specifies that, to prove that an event does not happen in the target, it is not enough to prove that it does not happen in the source: it is also necessary to prove that the source program is safe, i.e. it does not have any undefined behavior (second disjunct). Indeed, if it had an undefined behavior, its continuations could exhibit the unwanted event, and so would not be in the property. ⊡

Note that this relation can easily be generalized to other settings. For instance, consider the setting in which we reached a low-level language like machine code. Target traces can contain new events that cannot occur in the source: indeed, in modern architectures like x86 a compiler typically uses only a fraction of the available instruction set. Some instructions might even perform dangerous operations, such as writing to the hard drive. Formally, the source and target do not have the same events anymore.

Thus, we consider a source alphabet $\Sigma_{\mathsf{S}} = \Sigma \cup \{\mathsf{Goes_wrong}\}$, and a target alphabet $\boldsymbol{\Sigma_{\mathbf{T}}} = \Sigma \cup \Sigma'$. The trace relation is defined in the same way and we obtain the same

property mappings as above, except that since target traces now have more events (some of which may be dangerous), the arbitrary continuations of target traces get more interesting. For instance, consider a new event that represents writing data on the hard drive, and suppose we want to prove that this event cannot happen for a compiled program. Then, proving this property requires exactly proving that the source program exhibits no undefined behavior. More generally, what one can prove about target-only events can only be either that they cannot appear (because there is no undefined behavior) or that any of them can appear (in the case of undefined behavior).

3.2 Resource Exhaustion

In this section, we return to the discussion about resource exhaustion from §1.

Example 3.2 (Resource Exhaustion). We consider traces made of events drawn from Σ_S in the source, and $\mathbf{\Sigma_T} = \Sigma_S \cup \{\mathbf{Resource_Limit_Hit}\}$ in the target. Recall the trace relation for resource exhaustion:

$$\mathsf{s} \sim \mathbf{t} \iff \mathsf{s} = \mathbf{t} \vee \exists m \leq \mathsf{s}.\ \mathbf{t} = m \cdot \mathbf{Resource_Limit_Hit}$$

Formally, this relation is similar to the one for undefined behavior, except this time it is the target trace that is allowed to end early instead of the source trace.

By inspection of the definitions, we can see that the induced trace property mappings $\tilde{\sigma}$ and $\tilde{\tau}$ are equal to the following (🐓):

$$\tilde{\sigma}(\boldsymbol{\pi_T}) = \{\mathsf{s} \mid \mathsf{s} \in \boldsymbol{\pi_T}\} \cap \{\mathsf{s} \mid \forall m \leq \mathsf{s}.\ m \cdot \mathbf{Resource_Limit_Hit} \in \boldsymbol{\pi_T}\}$$

$$\tilde{\tau}(\pi_S) = \pi_S \cup \{m \cdot \mathbf{Resource_Limit_Hit} \mid \exists \mathsf{s} \in \pi_S.\ m \leq \mathsf{s}\}$$

These capture the following intuitions:

- The target-to-source mapping $\tilde{\sigma}$ states that, to prove a property of the compiled program, one has to prove that the traces of the source program satisfy two conditions: (1) they must also satisfy the target property; and (2) the termination of every one of their prefixes by a resource exhaustion error must satisfy the target property.
- The other mapping $\tilde{\tau}$ states that a compiled program produces traces that either belong to the same properties as the traces of the source program or end early due to resource exhaustion.

Note that $\tilde{\sigma}$ is rather restrictive: any property that prevents resource exhaustion cannot be proved using source-level reasoning. Indeed, if $\boldsymbol{\pi_T}$ does not allow resource exhaustion, then $\tilde{\sigma}(\boldsymbol{\pi_T}) = \varnothing$. This is to be expected since resource exhaustion is simply not accounted for at the source level.

On the other hand, in this model, safety properties [25] are mapped (in both directions) to other safety properties (🐓). This is a desirable trait for a relation: since safety properties are usually easier to reason about, one who is only interested in target safety properties can reason about them using powerful source-level tools for safety properties.

CakeML's compiler correctness theorem is an instance of $\mathsf{CC}^{\sim}$ for the $\sim$ relation above. We have also implemented two small compilers that are correct for this relation as illustrations of other situations where this relation is needed. The full details can be found in the Coq development in the supplementary materials.The first compiler (🐓) goes from a simple expression language, similar to the one in §3.3 but without inputs, to the same language where the execution is bounded by fuel: each execution

step consumes some amount of fuel and the execution immediately halts when it runs out of fuel. The compiler is the identity.

The second compiler (🐓) is more interesting: we proved this $\mathsf{CC}^{\sim}$ instance for a variant of Xavier Leroy's DSSS'17 compiler from a WHILE language to a simple stack machine [28]. We enriched the two languages with outputs and modified the semantics of the stack machine so that it falls into an error state if the stack reaches a certain size. The proof uses a standard forward simulation modified to account for failure. ⊡

Finally, we conclude this section by noting that the resource exhaustion relation and the undefined behavior relation from the previous subsection can easily be combined together. Indeed, given a relation $\sim_{\mathrm{UB}}$ and a relation $\sim_{\mathrm{RE}}$ defined as above on the same sets of traces, one can build a new relation $\sim$ that allows both refinement of undefined behavior and resource exhaustion by taking their union: $\mathsf{s} \sim \mathbf{t} \triangleq \mathsf{s} \sim_{\mathrm{UB}} \mathbf{t} \vee \mathsf{s} \sim_{\mathrm{RE}} \mathbf{t}$. A compiler that is $\mathsf{CC}^{\sim_{\mathrm{UB}}}$ or $\mathsf{CC}^{\sim_{\mathrm{RE}}}$ is trivially $\mathsf{CC}^{\sim}$, though the converse is not true.

3.3 Different Source and Target Values

We now illustrate trace-relating compilation for a compiler whose main feature is translating source-level booleans into target-level natural numbers. Given the simplicity of the languages and of the compiler, most of the details of the formalization are deferred to the online appendix.

The source language is a pure, statically typed expression language whose expressions e include naturals n, booleans b, conditionals, arithmetic and relational operations, boolean inputs $\mathsf{in_b}$ and natural inputs $\mathsf{in_n}$. A trace s is a list of inputs is paired with a result r, which can be a natural, a boolean or an error. Well-typed programs never produce error (🐓). Types ty are either N (naturals) or B (booleans); typing is standard. The source language has a big-step operational semantics ($\mathsf{e} \rightsquigarrow \langle \mathsf{is}, \mathsf{r} \rangle$) which tells how an expression e generates a trace $\langle \mathsf{is}, \mathsf{r} \rangle$. The semantics is standard and therefore omitted.

The target language is analogous to the source except that it is untyped, only has naturals $\mathbf{n}$ and its only inputs are naturals $\mathbf{in_n}$. The semantics of the target language is also given in big-step style. Since we only have naturals and all expressions operate on them, no error result is possible.

The compiler is homomorphic, translating a source expression to the same target expression; the only differences are natural numbers (and conditionals), as noted below.

$$\mathsf{true}{\downarrow} = \mathbf{1} \quad \mathsf{in_b}{\downarrow} = \mathbf{in_n} \qquad \mathsf{e_1} \leq \mathsf{e_2}{\downarrow} = \mathbf{if}\ \mathsf{e_1}{\downarrow} \leq \mathsf{e_2}{\downarrow}\ \mathbf{then}\ \mathbf{1}\ \mathbf{else}\ \mathbf{0}$$

$$\mathsf{false}{\downarrow} = \mathbf{0} \quad \mathsf{in_n}{\downarrow} = \mathbf{in_n} \quad \mathsf{if}\ \mathsf{e_1}\ \mathsf{then}\ \mathsf{e_2}\ \mathsf{else}\ \mathsf{e_3}{\downarrow} = \mathbf{if}\ \mathsf{e_1}{\downarrow} \leq \mathbf{0}\ \mathbf{then}\ \mathsf{e_3}{\downarrow}\ \mathbf{else}\ \mathsf{e_2}{\downarrow}$$

Compiling the *if-then-else* statement is more complicated: the target condition $\mathsf{e_1}{\downarrow} \leq \mathbf{0}$ is used to check that $\mathsf{e_1}$ is false, and therefore the *then* and *else* branches of the source are swapped in the target.

Relating Traces We relate basic values (naturals and booleans) in a non-injective fashion as noted below. Then, we extend the relation to lists of inputs in a pointwise fashion (Rules Empty and Cons) and lift that relation to traces (Rules Nat and Bool).

$$\mathsf{n} \sim \mathbf{n} \qquad \mathsf{true} \sim \mathbf{n} \quad \text{if } \mathbf{n} > 0 \qquad \mathsf{false} \sim \mathbf{0}$$

$$\frac{}{\varnothing \sim \varnothing}\ (\text{Empty}) \qquad \frac{\mathsf{i} \sim \mathbf{i} \quad \mathsf{is} \sim \mathbf{is}}{\mathsf{i} \cdot \mathsf{is} \sim \mathbf{i} \cdot \mathbf{is}}\ (\text{Cons}) \quad \Bigg| \quad \frac{\mathsf{is} \sim \mathbf{is} \quad \mathsf{n} \sim \mathbf{n}}{\langle \mathsf{is}, \mathsf{n} \rangle \sim \langle \mathbf{is}, \mathbf{n} \rangle}\ (\text{Nat}) \qquad \frac{\mathsf{is} \sim \mathbf{is} \quad \mathsf{b} \sim \mathbf{n}}{\langle \mathsf{is}, \mathsf{b} \rangle \sim \langle \mathbf{is}, \mathbf{n} \rangle}\ (\text{Bool})$$

With this relation, the compiler is proven correct in the sense of $\mathsf{CC}^{\sim}$.

Theorem 3.3 **($\cdot\!\downarrow$ is correct 🐓).** $\cdot\!\downarrow$ *is* $\mathsf{CC}^{\sim}$.

A notable difficulty in the proof of Theorem 3.3 arises from the trace-relating compilation setting: that proof does not follow from determinacy, input totality and forward simulation, as for compilation chains that have the same set of traces at both source and target level [7]. Instead, because the trace relation is not injective (both 0 and false are mapped to **0**), the type system is used to disambiguate between the two possibilities to back-translate a target trace. The property mappings $\tilde{\sigma}$ and $\tilde{\tau}$ induced by the trace relation $\sim$ defined above capture the intuition behind encoding booleans as naturals:

- the source-to-target mapping allows true to be encoded by any non-zero number;
- the target-to-source mapping requires that **0** is replaceable by, e.g., *both* 0 and false.

An extension to the property mappings is to take into account the *type* of the source program and obtain a type-aware property mapping. That is, in order to study the source obligation for source programs e of type ty, one can instead consider the following restriction: $\tilde{\sigma}_{\mathsf{ty}}(\boldsymbol{\pi_T}) = \tilde{\sigma}(\boldsymbol{\pi_T}) \cap \{\mathsf{s} \mid \exists \mathsf{e}' : \mathsf{ty}.\ \mathsf{e}' \rightsquigarrow \mathsf{s}\}$. This strengthens the proof obligation by restricting it to traces that can be produced by programs of a fixed type. Reasoning using this proof obligation is indeed correct, because when σ is used in TP^{σ}, it is always restricted to source traces produced by a program of the appropriate type. This narrowing of the source obligation can be used to guide verification: if $\tilde{\sigma}$ states that one has to prove that the result is either 0 or false, then $\tilde{\sigma}_{\mathsf{ty}}$ tells us which disjunct to prove: if the source programs have type N we need 0, otherwise false.

3.4 Abstraction Mismatches

We now consider how to relate traces where a single source action is compiled to multiple target ones. To illustrate this, we take a pure, statically typed source language that can output pairs of arbitrary size, and a pure, *untyped* target language where sent values have a fixed size. Concretely, the source is analogous to the language of §3.3 except that it does not have inputs or booleans and it has an expression send e, which can send a (nested) pair e of values in a single action. That is, given that e reduces to a pair, e.g., $\langle \mathsf{v1}, \langle \mathsf{v2}, \mathsf{v3} \rangle\rangle$, expression send $\langle \mathsf{v1}, \langle \mathsf{v2}, \mathsf{v3} \rangle\rangle$ emits action $\langle \mathsf{v1}, \langle \mathsf{v2}, \mathsf{v3} \rangle\rangle$.

That expression is compiled into a sequence of individual sends in the target language **send v1 ; send v2 ; send v3**, since in the target, **send e** sends the value that **e** reduces to, but the language has no pairs.

For space constraints we omit the full formalisation of these simple languages and of the homomorphic compiler ($(\cdot)\!\downarrow : \mathsf{e} \to \mathbf{e}$). The only interesting bit is the compilation of the send expression, which relies on the gensend $(\cdot)$ function below. That function takes a source expression of a given type and returns a sequence of target **send** instructions that send each element of the expression.

$$\texttt{gensend}\,(\vdash \mathsf{e} : \tau) = \begin{cases} \mathbf{send}\ (\vdash \mathsf{e} : \mathsf{N})\!\downarrow & \text{if } \tau = \mathsf{N} \\ \texttt{gensend}\,(\vdash \mathsf{e.1} : \tau')\,;\,\texttt{gensend}\,(\vdash \mathsf{e.2} : \tau'') & \text{if } \tau = \tau' \times \tau'' \end{cases}$$

Relating Traces We start with the trivial relation between numbers: $\mathsf{n} \sim^0 \mathbf{n}$, i.e. numbers are related when they are the same number. We cannot build a relation between single actions since a single source action is related to multiple target ones. So, we define a relation between a source action M and a target trace $\mathbf{t}$ inductively on the structure of M (which is a pair of values, and values are natural numbers or pairs).

$$\frac{\mathsf{n} \sim^0 \mathbf{n} \quad \mathsf{n}' \sim^0 \mathbf{n}'}{\langle \mathsf{n}, \mathsf{n}' \rangle \sim \mathbf{n} \cdot \mathbf{n}'}\ (\text{Trace-Rel-N-N}) \qquad \frac{\mathsf{n} \sim^0 \mathbf{n} \quad \mathsf{M} \sim \mathbf{t}}{\langle \mathsf{n}, \mathsf{M} \rangle \sim \mathbf{n} \cdot \mathbf{t}}\ (\text{Trace-Rel-N-M}) \qquad \frac{\mathsf{M} \sim \mathbf{t} \quad \mathsf{n} \sim^0 \mathbf{n}}{\langle \mathsf{M}, \mathsf{n} \rangle \sim \mathbf{t} \cdot \mathbf{n}}\ (\text{Trace-Rel-M-N}) \qquad \frac{\mathsf{M} \sim \mathbf{t} \quad \mathsf{M}' \sim \mathbf{t}'}{\langle \mathsf{M}, \mathsf{M}' \rangle \sim \mathbf{t} \cdot \mathbf{t}'}\ (\text{Trace-Rel-M-M})$$

A pair of naturals is related to the two actions that send each element of the pair (Rule Trace-Rel-N-N). If a pair is made of sub-pairs, we request all such sub-pairs to be related (Rules Trace-Rel-N-M to Trace-Rel-M-M).

We build on these rules to define the relation between source and target traces $\mathsf{s} \sim \mathbf{t}$ for which the compiler is correct (Theorem 3.4). Trivially, traces are related when they are both empty. Alternatively, given related traces, we can concatenate a source action and a second target trace provided that they are related (Rule Trace-Rel-Single). This relation induces the standard mappings between source and target properties, $\tilde{\tau}$ and $\tilde{\sigma}$ (Definition 2.5).

$$\frac{\mathsf{s} \sim \mathbf{t} \quad \mathsf{M} \sim \mathbf{t}'}{\mathsf{s} \cdot \mathsf{M} \sim \mathbf{t} \cdot \mathbf{t}'}\ (\text{Trace-Rel-Single})$$

Theorem 3.4 ($(\cdot)\downarrow$ is correct). $(\cdot)\downarrow$ *is* $\mathsf{CC}^{\sim}$.

With our trace relation, the trace property mappings capture the following intuitions:

- The target-to-source mapping states that a source property can reconstruct target action as it sees fit. For example, trace $\mathbf{4} \cdot \mathbf{6} \cdot \mathbf{5} \cdot \mathbf{7}$ is related to $\langle 4, 6 \rangle \cdot \langle 5, 7 \rangle$ and $\langle\langle 4, \langle 6, \langle 5, 7 \rangle\rangle\rangle\rangle$ (and many more variations). This gives freedom to the source implementation of a target behavior, which follows from the non-injectivity of $\sim$.[4]
- The source-to-target mapping states that the only valid target traces are those that source programs can compute, without any addition. This intuitively seems to preserve the meaning of the property, mapping source safety properties to target safety properties and even source hyperproperties to target hyperproperties.

4 Trace-Relating Compilation and Noninterference Preservation

When source and target observations are drawn from the same set, a correct compiler ($\mathsf{CC}^{=}$) is enough to ensure the preservation of all subset-closed hyperproperties (§2.5), in particular of *noninterference* (NI) [16]. This is possible only under the assumption that source and target observable actions are the same. In the scenario where target observations are strictly more informative than source observations, the best guarantee one may expect from a correct trace-relating compiler ($\mathsf{CC}^{\sim}$) is a *weakening* (or *declassification*) of target noninterference that matches the noninterference property satisfied in the source. To formalize this reasoning, this section applies the trinitarian view of trace-relating compilation to the general framework of abstract noninterference (ANI) [15].

We first define NI and explain the issue of preserving source NI via a $\mathsf{CC}^{\sim}$ compiler. We then introduce ANI, which allows characterizing various forms of noninterference,

[4] Making $\sim$ injective is a matter of adding open and close parentheses actions in target traces.

and formulate a general theory of ANI-preservation via $\mathsf{CC}^{\sim}$. We also study how to deal with cases such as undefined behavior in the target. Finally, we answer the dual question, i.e. how to recover source NI from target NI.

Intuitively, NI requires that publicly observable outputs do not reveal information about private inputs. To define this formally, we need a few additions to our setup. We indicate the (disjoint) *input* and *output* projections of a trace t as t° and $t^{\bullet}$ respectively. Denote with $[t]_{low}$ the equivalence class of a trace t, which is obtained using a standard low-equivalence relation that relates any two high (private) events, and that relates low (public) events only if they are equal. Then, NI for source traces can be defined as:

$$\mathsf{NI_S} = \{\pi_\mathsf{S} \mid \forall \mathsf{s}_1 \mathsf{s}_2 \in \pi_\mathsf{S}.\ [\mathsf{s}_1^{\circ}]_{low} = [\mathsf{s}_2^{\circ}]_{low} \Rightarrow [\mathsf{s}_1^{\bullet}]_{low} = [\mathsf{s}_2^{\bullet}]_{low}\}$$

In words, source NI comprises the sets of traces that have equivalent low output projections as soon as their low input projections are equivalent.

Trace-Relating Compilation and Noninterference. In a compilation chain where additional observations are possible in the target, it is a priori unclear whether a noninterfering source program is compiled to a noninterfering target program or not, and if so, whether the notion of NI in the target is the expected or desired one.

To illustrate this issue, consider a scenario where target traces additionally expose the execution time of a trace. Let $\mathsf{Trace_S}$ denote the set of traces in the source, and $\mathbf{Trace_T} = \mathsf{Trace_S} \times \mathbb{N}^{\omega}$ be the set of target traces, where $\mathbb{N}^{\omega} \triangleq \mathbb{N} \cup \{\omega\}$. Target traces have two components: a source trace, and a natural number that denotes the time spent to produce the trace (ω if infinite). We define the trace relation as:

$$\mathsf{s} \sim \mathbf{t} \iff \exists \mathbf{n}.\ \mathbf{t} = (\mathsf{s}, \mathbf{n})$$

Intuitively, a source trace is related to any target trace with the same first component, irrespective of the execution time component. Thus, a compiler is $\mathsf{CC}^{\sim}$ if any trace that can be exhibited in the target can be simulated in the source in any time. Building on this, the strongest notion of NI one can expect at the target level, $\mathbf{NI_T}$, is that private inputs do not affect public outputs (observable also in the source) or the execution time (observable only in the target). All properties $\boldsymbol{\pi}_\mathbf{T} \in \mathbf{NI_T}$ are obtained by some $\pi_\mathsf{S} \in \mathsf{NI_S}$, $\boldsymbol{n} \in \mathbb{N}^{\omega}$ and $\mathcal{I} \subseteq \mathbb{N}^{\omega}$ as follows. For every $\mathsf{s} \in \pi_\mathsf{S}$, if π_S contains some other trace that is low-equivalent on inputs to s, then $\boldsymbol{\pi}_\mathbf{T}$ contains s paired only with $\boldsymbol{n}$ and no other timing. If there are no such traces in π_S, then $\boldsymbol{\pi}_\mathbf{T}$ includes the whole $\{(\mathbf{s}, \boldsymbol{h}) \mid \boldsymbol{h} \in \mathcal{I}\}$. Applying the trinitarian view for subset-closed hyperproperties (§2.5) to NI yields that for a $\mathsf{CC}^{\sim}$ compilation chain, if W satisfies $\mathsf{NI_S}$, then $\mathsf{W}{\downarrow}$ satisfies $Cl_{\subseteq} \circ \tilde{\tau}(\mathsf{NI_S})$. However, can this resulting hyperproperty be interpreted as plain noninterference? Unfortunately not, as this is visibly a strictly weaker hyperproperty than $\mathbf{NI_T}$ since:

$$Cl_{\subseteq} \circ \tilde{\tau}(\mathsf{NI_S}) = Cl_{\subseteq}(\{\pi_\mathsf{S} \times \mathbb{N}^{\omega} \mid \pi_\mathsf{S} \in \mathsf{NI_S}\}) = \{\pi_\mathsf{S} \times \mathcal{I} \mid \pi_\mathsf{S} \in \mathsf{NI_S} \wedge \mathcal{I} \subseteq \mathbb{N}^{\omega}\},$$

the first equality coming from $\tilde{\tau}(\pi_\mathsf{S}) = \pi_\mathsf{S} \times \mathbb{N}^{\omega}$, and the second by $\mathsf{NI_S}$ being subset-closed. As we will see, this hyperproperty *can* be characterized as a form of NI, which one might call *timing-insensitive noninterference*, which is ensured only against attackers that cannot measure execution time. For this characterization, and to describe different forms of noninterference as well as formally analyze their preservation by a $\mathsf{CC}^{\sim}$ compiler, we rely on the general framework of *abstract noninterference* [15].

Abstract Noninterference. ANI is a generalization of NI whose formulation relies on properties (i.e., abstractions in abstract interpretation [13]) in order to encompass arbitrary variants of NI. ANI is parameterized by an *observer abstraction* ρ, which denotes the distinguishing power of the attacker, and a *selection abstraction* ϕ, which specifies when to check NI, and therefore captures a form of declassification [41].[5] Formally:

$$ANI^{\rho}_{\phi} = \{\pi \mid \forall t_1 t_2 \in \pi.\ \phi(t_1^{\circ}) = \phi(t_2^{\circ}) \Rightarrow \rho(t_1^{\bullet}) = \rho(t_2^{\bullet})\}$$

By picking $\phi = \rho = [\cdot]_{low}$, we recover the standard noninterference defined above, where NI must hold for all low inputs (i.e., no declassification of private inputs), and the observational power of the attacker is limited to distinguishing low outputs.

Moreover, the observational power of the attacker can be weakened by choosing a more liberal relation for ρ. For instance, one may limit the attacker to observe the *parity* of output integer values. Another way to weaken ANI is to use ϕ to specify that noninterference is only required to hold for a subset of low inputs.

To be formally precise, ϕ and ρ are defined over sets of (input and output projections of) traces, so when we write $\phi(t)$ above, this should be understood as a convenience notation for $\phi(\{t\})$. Likewise, $\phi = [\cdot]_{low}$ should be understood as $\phi = \lambda\pi. \bigcup_{t\in\pi}[t]_{low}$, i.e., the powerset lifting of $[\cdot]_{low}$. Additionally, ϕ and ρ are required to be upper-closed operators (uco)—i.e., monotonic, idempotent and extensive—on the poset that is the powerset of traces ordered by inclusion [15].

Trace-Relating Compilation and ANI for Timing. We can now reformulate our example with execution time observable in the target in terms of ANI. We have $\mathsf{NI_S} = ANI^{\phi}_{\rho}$ with $\phi = \rho = [\cdot]_{low}$. In this case, we can formally describe the hyperproperty that a compiled program $\mathsf{W}\!\downarrow$ satisfies whenever W satisfies $\mathsf{NI_S}$ as an instance of ANI:

$$Cl_{\subseteq} \circ \tilde{\tau}(\mathsf{NI_S}) = ANI^{\boldsymbol{\rho}}_{\boldsymbol{\phi}}$$

$$\text{for } \boldsymbol{\phi} = \phi \text{ and } \boldsymbol{\rho}(\boldsymbol{\pi_T}) = \{(\mathsf{s}, \mathbf{n}) \mid \exists(\mathsf{s_1}, \mathbf{n_1}) \in \boldsymbol{\pi_T}.\ [\mathsf{s}^{\bullet}]_{low} = [\mathsf{s}_1^{\bullet}]_{low}\}$$

The definition of $\boldsymbol{\phi}$ tells us that the trace relation does not affect the selection abstraction. The definition of $\boldsymbol{\rho}$ characterizes an observer that cannot distinguish execution times for noninterfering traces (notice that $\mathbf{n_1}$ in the definition of $\boldsymbol{\rho}$ is discarded). For instance, $\boldsymbol{\rho}(\{(\mathsf{s}, \mathbf{n_1})\}) = \boldsymbol{\rho}(\{(\mathsf{s}, \mathbf{n_2})\})$, for any s, $\mathbf{n_1}$, $\mathbf{n_2}$. Therefore, in this setting, we know explicitly through $\boldsymbol{\rho}$ that a $\mathsf{CC}^{\sim}$ compiler "degrades" source noninterference to target *timing-insensitive* noninterference.

Trace-Relating Compilation and ANI in General. While the definitions of $\boldsymbol{\phi}$ and $\boldsymbol{\rho}$ above can be discovered by intuition, we want to know whether there is a systematic way of obtaining them. In other words, for *any* trace relation $\sim$ and *any* notion of source NI, what property is guaranteed on noninterfering source programs by a $\mathsf{CC}^{\sim}$ compiler?

We can now answer this question generally (Theorem 4.1): any source notion of noninterference expressible as an instance of ANI is mapped to a corresponding notion of noninterference in the target, whenever source traces are an abstraction of target ones (i.e., when $\sim$ is a total and surjective map).

So we now consider trace relations that can be split into input and output trace relations (denoted as $\sim \triangleq \langle \overset{\circ}{\sim}, \overset{\bullet}{\sim} \rangle$) such that $\mathsf{s} \sim \mathbf{t} \iff \mathsf{s}^{\circ} \overset{\circ}{\sim} \mathbf{t}^{\circ} \wedge \mathsf{s}^{\bullet} \overset{\bullet}{\sim} \mathbf{t}^{\bullet}$. The trace

[5] ANI includes a third parameter η, which describes the maximal input variation that the attacker may control. Here we omit η (i.e., take it to be the identity) in order to simplify the presentation.

relation $\sim$ induces a Galois connection between the sets of trace properties $\tilde{\tau} \leftrightarrows \tilde{\sigma}$ as described in §2.3. Similarly, both $\overset{\circ}{\sim}$ and $\overset{\bullet}{\sim}$ induce Galois connections, $\tilde{\tau}^\circ \leftrightarrows \tilde{\sigma}^\circ$ and $\tilde{\tau}^\bullet \leftrightarrows \tilde{\sigma}^\bullet$ between the sets of input and output properties. In the timing example, time is an output so we have $\sim \triangleq \langle =, \overset{\bullet}{\sim} \rangle$ and $\overset{\bullet}{\sim}$ is defined as $\mathsf{s}^\bullet \overset{\bullet}{\sim} \mathbf{t}^\bullet \iff \exists \mathbf{n}.\ \mathbf{t}^\bullet = (\mathsf{s}^\bullet, \mathbf{n})$.

Theorem 4.1 (Compiling ANI). *Assume traces of source and target languages are related via $\sim \subseteq \mathsf{Trace_S} \times \mathbf{Trace_T}$, $\sim \triangleq \langle \overset{\circ}{\sim}, \overset{\bullet}{\sim} \rangle$ such that $\overset{\circ}{\sim}$ and $\overset{\bullet}{\sim}$ are both total, surjective maps from target to source traces. Assume $\downarrow$ is a $\mathsf{CC}^\sim$ compiler, and $\phi \in uco(2^{\mathsf{Trace}^\circ_\mathsf{S}})$, $\rho \in uco(2^{\mathsf{Trace}^\bullet_\mathsf{S}})$.*

If W satisfies ANI^{ρ}_{ϕ}, then $\mathsf{W}\!\downarrow$ satisfies $ANI^{\rho^\#}_{\phi^\#}$ where $\phi^\#$ and $\rho^\#$ are defined as:

$$\phi^\# = g^\circ \circ \phi \circ f^\circ \qquad\qquad \rho^\# = g^\bullet \circ \rho \circ f^\bullet \quad \textit{where}$$

$$f^\circ(\pi^\circ_\mathbf{T}) = \left\{ \mathsf{s}^\circ \;\middle|\; \exists \mathbf{t}^\circ \in \pi^\circ_\mathbf{T}.\ \mathsf{s}^\circ \overset{\circ}{\sim} \mathbf{t}^\circ \right\} \qquad g^\circ(\pi^\circ_\mathsf{S}) = \{ \mathbf{t}^\circ \mid \forall \mathsf{s}^\circ.\ \mathsf{s}^\circ \overset{\circ}{\sim} \mathbf{t}^\circ \Rightarrow \mathsf{s}^\circ \in \pi^\circ_\mathsf{S} \}$$

(and both $f^\bullet$ and $g^\bullet$ are defined similarly)

As expected, $\phi^\# = g^\circ \circ \phi \circ f^\circ = \phi$ and $\rho^\# = g^\bullet \circ \rho \circ f^\bullet = \rho$, recovering the definitions we justified intuitively above for timing. Moreover, we can prove that in general $ANI^{\rho^\#}_{\phi^\#} \subseteq Cl_\subseteq \circ \tilde{\tau}(ANI^{\rho}_{\phi})$. Therefore, the derived guarantee $ANI^{\rho^\#}_{\phi^\#}$ is at least as strong as the one that follows by knowing that the compiler $\downarrow$ is $\mathsf{CC}^\sim$.

Noninterference and Undefined Behavior. In its statement, Theorem 4.1 does not apply to several scenarios from §3 such as undefined behavior (§3.1), because in those cases the relation $\overset{\bullet}{\sim}$ is not a total surjective map. Nevertheless, one can still exploit our framework to reason about the impact of compilation on noninterference.

Let us consider $\sim \triangleq \langle \overset{\circ}{\sim}, \overset{\bullet}{\sim} \rangle$ where $\overset{\circ}{\sim}$ is any total and surjective map from target to source inputs (e.g., equality) and $\overset{\bullet}{\sim}$ is defined as $\mathsf{s}^\bullet \overset{\bullet}{\sim} \mathbf{t}^\bullet \iff \mathsf{s}^\bullet = \mathbf{t}^\bullet \vee \exists m^\bullet \leq \mathbf{t}^\bullet.\ \mathsf{s}^\bullet = m^\bullet \cdot Goes_wrong$. Intuitively, a $\mathsf{CC}^\sim$ compiler guarantees that any interference cannot be observed by a target attacker that cannot exploit undefined behavior to learn private information. This intuition can be made formal by the following theorem.

Theorem 4.2 (Relaxed Compiling ANI). *Relax the assumptions of Theorem 4.1 by allowing $\overset{\bullet}{\sim}$ to be* any *output trace relation. If W satisfies ANI^{ρ}_{ϕ}, then $\mathsf{W}\!\downarrow$ satisfies $ANI^{\rho^\#}_{\phi^\#}$ where $\phi^\#$ is defined as in Theorem 4.1, and $\rho^\#$ is such that:*

$$\forall \mathsf{s}\ \mathbf{t}.\ \mathsf{s}^\bullet \overset{\bullet}{\sim} \mathbf{t}^\bullet \Rightarrow \rho^\#(\mathbf{t}^\bullet) = \rho^\#(\tilde{\tau}^\bullet(\rho(\mathsf{s}^\bullet)))$$

Technically, instead of giving us a *definition* of $\rho^\#$, the theorem gives a *property* of it. The property states that, given a target output trace $\mathbf{t}^\bullet$, the attacker cannot distinguish it from any other target output traces produced by other possible compilations ($\tilde{\tau}^\bullet$) of the source trace s it relates to, up to the observational power of the source level attacker ρ. Therefore, given a source attacker ρ, the theorem characterizes a *family* of attackers that cannot observe any interference for a correctly compiled noninterfering program. Notice that the target attacker $\rho^\# = \lambda_.\ \top$ satisfies the premise of the theorem, but defines a trivial hyperproperty, so that we cannot prove in general that $ANI^{\rho^\#}_{\phi^\#} \subseteq Cl_\subseteq \circ \tilde{\tau}(ANI^{\rho}_{\phi})$. The same $\rho^\# = \lambda_.\ \top$, shows that the family of attacker described in Theorem 4.2 is non empty, and this ensures the existence of the most powerful of them [15], whose explicit characterization we leave as future work.

From Target NI to Source NI. We now explore the dual question: under which hypotheses does trace-relating compiler correctness alone reduce satisfaction of target noninterference to source noninterference? This question is of practical interest, as one would be allowed to protect from target attackers by ensuring noninterference in the source, usually an easier task if the source language has some enforcement mechanism such as a security type system [1].

Let us consider an extension of the languages from §3.4 with inputting (pairs of) values as well. It is easy to show that the compiler described in §3.4 is still $\mathsf{CC}^{\sim}$. Assume that we want to satisfy a given notion of target noninterference after compilation, i.e. $\mathsf{W}{\downarrow} \models ANI^{\rho}_{\phi}$. Recall that the observational power of the target attacker, ρ, is expressed as a property of sequences of values. To express the same property, or attacker, in the source, we have to abstract the way pairs of values are nested. For instance, the source attacker should not distinguish $\langle \mathsf{v}_1, \langle \mathsf{v}_2, \mathsf{v}_3 \rangle\rangle$ and $\langle\langle \mathsf{v}_1, \mathsf{v}_2 \rangle, \mathsf{v}_3\rangle$. In general (i.e. when $\mathring{\sim}$ is not the identity), this argument is valid only when ϕ can be represented in the source. More precisely, ϕ must consider as equivalent all target inputs that are related to the same source one, this because in the source it is not possible to have a finer distinction of inputs. This intuitive correspondence can be formalized as follows:

Theorem 4.3 (Target ANI by source ANI). *Let $\phi \in uco(2^{\mathbf{Trace}^{\circ}_{\mathrm{T}}})$, $\rho \in uco(2^{\mathbf{Trace}^{\bullet}_{\mathrm{T}}})$ and $\dot{\sim}$ a total and surjective map from source outputs to target ones and assume that*

$$\forall \mathsf{s}\ \mathbf{t}.\ \mathsf{s}^{\circ} \mathring{\sim} \mathbf{t}^{\circ} \Rightarrow \phi(\mathbf{t}^{\circ}) = \phi(\tilde{\tau}^{\circ}(\mathsf{s}^{\circ}))$$

If $\cdot\!\downarrow$ is a $\mathsf{CC}^{\sim}$ compilation chain and W satisfies $ANI^{\rho^{\#}}_{\phi^{\#}}$, then $\mathsf{W}{\downarrow}$ satisfies ANI^{ρ}_{ϕ} for

$$\phi^{\#} = \tilde{\sigma}^{\circ} \circ \phi \circ \tilde{\tau}^{\circ} \qquad\qquad \rho^{\#} = \tilde{\sigma}^{\bullet} \circ \rho \circ \tilde{\tau}^{\bullet}$$

To wrap up the discussion about noninterference, the results presented in this section formalize and generalize some intuitive facts about compiler correctness and noninterference. Of course, they all place some restrictions on the shape of the noninterference instances that can be considered, because compiler correctness alone is in general not a strong enough criterion for dealing with many security properties (see e.g., [6, 14]).

5 Trace-Relating Secure Compilation

So far we have studied compiler correctness criteria for whole, standalone programs. However, in practice, programs do not exist in isolation, but in a context where they interact with other programs, libraries, etc. In many cases, this context cannot be assumed to be benign and could instead behave maliciously to try to disrupt a compiled program.

Hence, in this section we consider the following *secure compilation* scenario: a source program is compiled and linked with an arbitrary target-level context, i.e. one that may not be expressible as the compilation of a source context. Compiler correctness does not address this case, as it does not consider arbitrary target contexts, looking instead at whole programs (empty context [27]) or well-behaved target contexts that behave like source ones (compositional compiler correctness [21, 24, 34, 44]).

To account for this scenario, Abate et al. [2] describe several secure compilation criteria based on the preservation of classes of (hyper)properties (e.g., trace properties, safety, hypersafety, hyperproperties, etc.) against arbitrary target contexts. For each of

these criteria, they give an equivalent "property-free" criterion, analogous to the equivalence between TP and $\mathsf{CC}^{=}$. For instance, their *robust* trace property preservation criterion (RTP) states that, for any trace property π, if program P plugged into any context C satisfies π, then the compiled program $\mathsf{P}\downarrow$ plugged into any target context $\mathbf{C}$ satisfies the π. Their equivalent criterion to RTP is RTC, which states that for any trace produced by the compiled program when linked with any target context, there is a source context that produces the same trace. Formally (writing $C\,[P]$ to mean program P linked with context C) they define:

$$\mathsf{RTP} \equiv \forall \pi.\ \forall \mathsf{P}.(\forall \mathsf{C}.\ \forall \mathsf{s}.\mathsf{C}\,[\mathsf{P}]\rightsquigarrow s \Rightarrow s \in \pi) \Rightarrow (\forall \mathbf{C}.\ \forall \mathbf{t}.\mathbf{C}\,[\mathsf{P}\downarrow]\rightsquigarrow t \Rightarrow t \in \pi)$$

$$\mathsf{RTC} \equiv \forall \mathsf{P}.\ \forall \mathbf{C}.\forall t.\mathbf{C}\,[\mathsf{P}\downarrow]\rightsquigarrow t \Rightarrow \exists \mathsf{C}.\ \mathsf{C}\,[\mathsf{P}]\rightsquigarrow t$$

In the following we adopt the notation $P \models_R \pi$ to mean "P *robustly* satisfies π" and write RTP more compactly as $\mathsf{RTP} \equiv \forall \pi.\ \forall \mathsf{P}.\ \mathsf{P} \models_{\mathsf{R}} \pi \Rightarrow \mathsf{P}\downarrow \models_{\mathbf{R}} \pi$.

All the criteria of Abate et al. [2] are stated in a setting where source and target traces are the same. In this section, we extend their result to our trace-relating setting, obtaining trintarian views for secure compilation. Despite the similarities with §2, more challenges show up, in particular when considering the robust preservation of proper sub-classes of trace properties. For example, the application of $\tilde{\sigma}$ or $\tilde{\tau}$ may lose the information that a property is safety, a crucial point for the equivalence with the property-free criterion for safety properties by Abate et al. [2]. We solve this problem by interpreting the class of safety properties as an *abstraction* of the class of all trace properties induced by a closure operator (§5.1). The rest of the section provides instances satisfying some secure compilation trinities, namely for the class of trace properties (§5.2), and for the proper sub-classes of safety properties and hypersafety (§5.3).

5.1 Trace-Relating Secure Compilation: A Spectrum of Trinities

In this section we generalize many of the criteria of Abate et al. [2] using the ideas of §2. The ideas of §2 are enough to easily generalize the criteria for trace properties and subset-closed hyperproperties of Abate et al. [2]. Before discussing how we solve the additional challenges of classes such as safety and hypersafety, we first explain the simple generalization of RTC to the trace-relating setting ($\mathsf{RTC}^{\sim}$) and its corresponding trinitarian view (Theorem 5.1):

Theorem 5.1 (Trinity for Robust Trace Properties 🐓). *For any trace relation $\sim$ and induced property mappings $\tilde{\tau}$ and $\tilde{\sigma}$, we have:* $\mathsf{RTP}^{\tilde{\tau}} \iff \mathsf{RTC}^{\sim} \iff \mathsf{RTP}^{\tilde{\sigma}}$, *where*

$$\mathsf{RTC}^{\sim} \equiv \forall \mathsf{P}\ \forall \mathbf{C_T}\ \forall \mathbf{t}.\ \mathbf{C_T}\,[\mathsf{P}\downarrow]\rightsquigarrow \mathbf{t} \Rightarrow \exists \mathsf{C_S}\ \exists \mathsf{s}.\ \mathsf{C_S}\,[\mathsf{P}]\rightsquigarrow \mathsf{s} \wedge \mathsf{s} \sim \mathbf{t}$$

$$\mathsf{RTP}^{\tilde{\tau}} \equiv \forall \mathsf{P}\ \forall \pi_{\mathsf{S}} \in 2^{\mathsf{Trace_S}}.\ \mathsf{P} \models_{\mathsf{R}} \pi_{\mathsf{S}} \Rightarrow \mathsf{P}\downarrow \models_{\mathbf{R}} \tilde{\tau}(\pi_{\mathsf{S}})$$

$$\mathsf{RTP}^{\tilde{\sigma}} \equiv \forall \mathsf{P}\ \forall \boldsymbol{\pi_T} \in 2^{\mathbf{Trace_T}}.\ \mathsf{P} \models_{\mathsf{R}} \tilde{\sigma}(\boldsymbol{\pi_T}) \Rightarrow \mathsf{P}\downarrow \models_{\mathbf{R}} \boldsymbol{\pi_T}$$

Abate et al. [2] propose many more equivalent pairs of criteria, each preserving different classes of (hyper)properties, which we briefly recap now. For trace properties, they also have criteria that preserve safety properties plus their version of liveness properties. For hyperproperties, they have criteria that preserve hypersafety properties, subset-closed hyperproperties, and arbitrary hyperproperties. Finally, they define *relational* hyperproperties, which are relations between the behaviours of multiple pro-

grams for expressing, e.g., that a program always runs faster than another. For relational hyperproperties, they have criteria that preserve arbitrary relational properties, relational safety properties, relational hyperproperties and relational subset-closed hyperproperties. Roughly speaking, the security guarantees due to robust preservation of trace properties regard only protecting the integrity of the program from the context, the guarantees of hyperproperties also regard data confidentiality, and the guarantees of relational hyperproperties even regard code confidentiality. Naturally, these stronger guarantees are increasingly harder to enforce and prove.

While we have lifted the most significant criteria from Abate et al. [2] to our trinitarian view, due to space constraints we provide the formal definitions only for the two most interesting criteria. We collect the generalization of many other criteria in Figure 2, which we outline below. Omitted definitions are available in the online appendix.

Beyond traces properties: robust safety and hyperproperty preservation. We detail robust preservation of safety properties and of arbitrary hyperproperties since they are both relevant from a security point of view and their generalization is interesting.

Theorem 5.2 (Trinity for Robust Safety Properties 🐓). *For any trace relation $\sim$ and for the induced property mappings $\tilde{\tau}$ and $\tilde{\sigma}$, we have:*

$$\mathsf{RTP}^{\mathit{Safe}\circ\tilde{\tau}} \iff \mathsf{RSC}^{\sim} \iff \mathsf{RSP}^{\tilde{\sigma}} \qquad \textit{where}$$

$$\mathsf{RSC}^{\sim} \equiv \forall \mathsf{P}\ \forall \mathbf{C_T}\ \forall \mathbf{t}\ \forall \mathbf{m} \leq \mathbf{t}.\mathbf{C_T}\left[\mathsf{P}{\downarrow}\right]{\rightsquigarrow}\mathbf{t} \Rightarrow \exists \mathsf{C_S}\ \exists \mathbf{t'} \geq \mathbf{m}\ \exists \mathsf{s} \sim \mathbf{t'}.\ \mathsf{C_S}\left[\mathsf{P}\right]{\rightsquigarrow}\mathsf{s}$$

$$\mathsf{RTP}^{\mathit{Safe}\circ\tilde{\tau}} \equiv \forall \mathsf{P} \forall \pi_{\mathsf{S}} \in 2^{\mathsf{Trace_S}}.\mathsf{P} \models_{\mathsf{R}} \pi_{\mathsf{S}} \Rightarrow \mathsf{P}{\downarrow} \models_{\mathbf{R}} (\mathit{Safe}\circ\tilde{\tau})(\pi_{\mathsf{S}})$$

$$\mathsf{RSP}^{\tilde{\sigma}} \equiv \forall \mathsf{P} \forall \boldsymbol{\pi}_{\mathbf{T}} \in \mathbf{Safety_T}.\mathsf{P} \models_{\mathsf{R}} \tilde{\sigma}(\boldsymbol{\pi}_{\mathbf{T}}) \Rightarrow \mathsf{P}{\downarrow} \models_{\mathbf{R}} \boldsymbol{\pi}_{\mathbf{T}}$$

There is an interesting asymmetry between the last two characterizations above, that we explain now in more detail. $\mathsf{RSP}^{\tilde{\sigma}}$ quantifies over target safety properties, while $\mathsf{RTP}^{\mathit{Safe}\circ\tilde{\tau}}$ quantifies over *arbitrary* source properties, but imposes the composition of $\tilde{\tau}$ with Safe, which maps an arbitrary target property $\boldsymbol{\pi}_{\mathbf{T}}$ to the target safety property that best over-approximates $\boldsymbol{\pi}_{\mathbf{T}}$[6] (an analogous *closure* was needed for subset-closed hyperproperties in Theorem 2.12). More precisely Safe is a closure operator on target properties, with $\mathbf{Safety_T} = \left\{\mathit{Safe}(\boldsymbol{\pi}_{\mathbf{T}}) \;\middle|\; \boldsymbol{\pi}_{\mathbf{T}} \in 2^{\mathbf{Trace_T}}\right\}$. The mappings

$$\mathit{Safe}\circ\tilde{\tau} : 2^{\mathsf{Trace_S}} \leftrightarrows 2^{\mathbf{Trace_T}} : \tilde{\sigma}\circ\iota\circ\mathit{Safe}$$

where $\iota = id_{|\mathbf{Safety_T}}$, determine a new Galois connection between trace properties.

A "robust" variant (🐓) of Theorem 2.4 gives us

$$\mathsf{RTP}^{\mathit{Safe}\circ\tilde{\tau}} \iff \mathsf{RTP}^{\tilde{\sigma}\circ\iota\circ\mathit{Safe}}$$

Finally, note that quantifying over arbitrary target properties, closing them in the class of safety and then applying $\tilde{\sigma}$, is the same as restricting $\tilde{\sigma}$ only to the properties in $\mathbf{Safety_T}$, so that $\mathsf{RTP}^{\mathit{Safe}\circ\tilde{\tau}}$ becomes exactly $\mathsf{RSP}^{\tilde{\sigma}}$, and $\mathsf{RTP}^{\mathit{Safe}\circ\tilde{\tau}} \iff \mathsf{RSP}^{\tilde{\sigma}}$.

This argument generalizes to arbitrary closure operators on target properties and on hyperproperties as long as the corresponding class is a sub-class of subset-closed hyperproperties. A few more trinities in the diagram rely on instances of this argument. This argument explains all but one of the asymmetries in Figure 2, which regards the robust preservation of arbitrary hyperproperties:

[6] $\mathit{Safe}(\boldsymbol{\pi}_{\mathbf{T}}) = \cap\{\boldsymbol{S}_{\mathbf{T}} \mid \boldsymbol{\pi}_{\mathbf{T}} \subseteq \boldsymbol{S}_{\mathbf{T}} \wedge \boldsymbol{S}_{\mathbf{T}} \in \mathbf{Safety_T}\}$, is the topological closure in the topology of Clarkson and Schneider [11] where safety properties coincide with the closed sets.

Theorem 5.3 (Weak Trinity for Robust Hyperproperties). *For a trace relation* $\sim \subseteq \mathsf{Trace_S} \times \mathbf{Trace_T}$*, and induced property mappings* $\tilde{\sigma}$ *and* $\tilde{\tau}$*,* $\mathsf{RHC}^{\sim}$ *is equivalent to* $\mathsf{RHP}^{\tilde{\tau}}$ *below. Moreover, if* $\tilde{\tau} \leftrightarrows \tilde{\sigma}$ *is a Galois insertion (i.e.* $\tilde{\tau} \circ \tilde{\sigma} = id$*),* $\mathsf{RHC}^{\sim}$ *implies* $\mathsf{RHP}^{\tilde{\sigma}}$*, while if* $\tilde{\sigma} \leftrightarrows \tilde{\tau}$ *is a Galois reflection (i.e.* $\tilde{\sigma} \circ \tilde{\tau} = id$*),* $\mathsf{RHP}^{\tilde{\sigma}}$ *implies* $\mathsf{RHC}^{\sim}$*.*

$$\mathsf{RHC}^{\sim} \equiv \forall \mathsf{P}\, \forall \mathbf{C_T}\, \exists \mathsf{C_S}\, \forall \mathbf{t}.\ \mathbf{C_T}\, [\mathsf{P}{\downarrow}] \rightsquigarrow \mathbf{t} \iff (\exists \mathsf{s} \sim \mathbf{t}.\ \mathsf{C_S}\, [\mathsf{P}] \rightsquigarrow \mathsf{s})$$

$$\mathsf{RHP}^{\tilde{\tau}} \equiv \forall \mathsf{P}\, \forall \mathsf{H_S}.\ \mathsf{P} \models_{\mathsf{R}} \mathsf{H_S} \Rightarrow \mathsf{P}{\downarrow} \models_{\mathbf{R}} \tilde{\tau}(\mathsf{H_S})$$

$$\mathsf{RHP}^{\tilde{\sigma}} \equiv \forall \mathsf{P}\, \forall \mathbf{H_T}.\ \mathsf{P} \models_{\mathsf{R}} \tilde{\sigma}(\mathbf{H_T}) \Rightarrow \mathsf{P}{\downarrow} \models_{\mathbf{R}} \mathbf{H_T}$$

This trinity is *weak* since extra hypotheses are needed to prove an implication. While the equivalence $\mathsf{RHC}^{\sim} \iff \mathsf{RHP}^{\tilde{\tau}}$ holds unconditionally, the other two implications hold only under distinct, stronger assumptions. For $\mathsf{RHP}^{\tilde{\sigma}}$ it is still possible and correct to deduce a source obligation for a given target hyperproperty $\mathbf{H_T}$ when no information is lost in the the composition $\tilde{\tau} \circ \tilde{\sigma}$ (i.e. the two maps are a Galois *insertion*). On the other hand, $\mathsf{RHP}^{\tilde{\tau}}$ is a consequence of $\mathsf{RHP}^{\tilde{\sigma}}$ when no information is lost in composing in the other direction, $\tilde{\sigma} \circ \tilde{\tau}$ (i.e. the two maps are a Galois *reflection*).

Navigating the diagram. For a given trace relation $\sim$, Figure 2 orders the generalized criteria according to their relative strength. If a trinity implies another (denoted by $\Rightarrow$), then the former provides stronger security for a compilation chain than the latter.

As mentioned, some property-free criteria (i.e., those for subset-closed hyperproperties, safety, hypersafety, 2-relational safety and 2-relational hyperproperties) in the diagram are derived from a stronger criterion closed with an additional operator. We have already presented the $Safe$ operator; other operators are $Cl_{\subseteq}$, $HSafe$, and $2rSafe$, which respectively ensure that the function they are composed with maps a subset-closed hyperproperty (as in Theorem 2.12), a hypersafety property, or a 2-relational safety property to the same class across languages. As a reading aid, we use the same background color for both the trinity from which the stronger criterion is borrowed and for the criterion using the closing operator. The trace-based trinity is blue, the subset-closed hyperproperty one is red, and 2-relational property one is green.

We now describe how to interpret the acronyms in Figure 2. All criteria start with R meaning they refer to robust preservation. Criteria for relational hyperproperties—here only arity 2 is shown for simplicity—contain 2r. Next, criteria names spell the class of hyperproperties they preserve: H for hyperproperties, SCH for subset-closed hyperproperties, HS for hypersafety, T for trace properties and S for safety properties. Finally, property-free criteria end with a C while property-full criteria involving $\tilde{\sigma}$ and $\tilde{\tau}$ end with P. Thus, *robust (*R*) subset-closed hyperproperty-preserving (*SCH*) compilation (*C*)* is $\mathsf{RSCHC}^{\sim}$, *robust (*R*) two-relational (*2r*) safety-preserving (*S*) compilation (*C*)* is $\mathsf{R2rSC}^{\sim}$ etc.

5.2 Instance of Trace-Relating Robust Preservation of Trace Properties

This section illustrates trace-relating secure compilation by extending the example of §3.3 to a target language having strictly more events than the source.

Source and Target Languages. The source and target languages used in this section (R^{S} and $\mathbf{R}^{\mathbf{T}}$, respectively) are nearly identical expression languages borrowing from the syntax of the source language of §3.3. Both languages add *sequencing* of expressions,

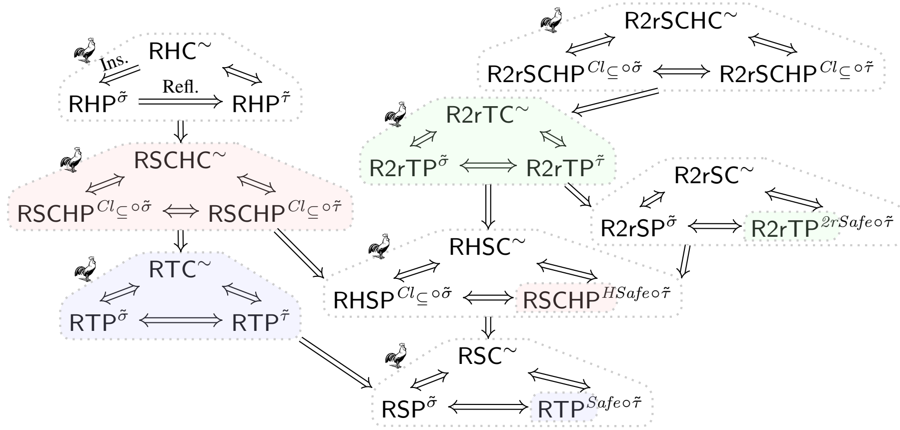

R robust	2r 2-relational	
H hyperproperties	SCH subset-closed hyperproperties	HS hypersafety
T trace properties	S safety properties	
P property-full criterion	C property-free criterion based on σ and τ	

Fig. 2: Hierarchy of trinitarian views of secure compilation criteria preserving classes of hyperproperties and the key to read each acronym. Shorthands 'Ins.' and 'Refl.' stand for Galois Insertion and Reflection. The 🐓 denotes trinities proven in Coq.

two kinds of *output events*, and the expressions that generate them: $\mathsf{out_S}\ \mathsf{n}$ and $\mathbf{out_S}\ \mathbf{n}$ (usable respectively in $\mathsf{R^S}$ and $\mathbf{R^T}$) and $\mathbf{out_T}\ \mathbf{n}$ (usable only in $\mathbf{R^T}$). The presence of $\mathbf{out_T}$ events is the sole difference between source and target. The extra events in the target model the fact that the target language has an increased ability to perform certain operations, some of them potentially dangerous (such as writing to the hard drive), which cannot be observed by the source language, and against which source-level reasoning can therefore offer no protection.

Both languages and compilation chains now deal with partial programs, contexts and linking of those two to produce whole programs. In this setting, a whole program is the combination of a *main expression* to be evaluated and a set of *function definitions* (with distinct names) that can refer to their argument symbolically and can be called by the main expression and by other functions. The set of functions of a whole program is the union of the functions of a partial program and a context; the latter also contains the main expression. The extensions of the typing rules and the operational semantics for whole programs are unsurprising and therefore elided. The trace model also follows closely that of §3.3: it consists of a list of *regular events* (including the new outputs) terminated by a *result event*. Finally, a partial program and a context can be linked into a whole program when their functions satisfy the requirements mentioned above.

The target language $\mathbf{R^T}$ is a small extension of the source, which allows target-only events $\mathbf{out_T}$ to appear in target expressions and consequently in target traces $\mathbf{t}$.

Relating Traces. In the present model, source and target traces differ only in the fact that the target draws (regular) events from a strictly larger set than the source, i.e. $\mathbf{\Sigma_T} \supset$

Σ_S. A natural relation between source and target traces essentially maps to a given target trace $\mathbf{t}$ the source trace that erases from $\mathbf{t}$ those events that exist only at the target level. Let $\mathbf{t}|_{\Sigma_S}$ indicate trace $\mathbf{t}$ filtered to retain only those elements included in alphabet Σ_S. We define the trace relation as:

$$s \sim \mathbf{t} \iff s = \mathbf{t}|_{\Sigma_S}$$

In the opposite direction, a source trace s is related to many target traces, as a finite number of target-only events can be inserted at any point in s. The induced mappings for this relation are:

$$\tilde{\tau}(\pi_S) = \{\mathbf{t} \mid \exists s.\, s = \mathbf{t}|_{\Sigma_S} \wedge s \in \pi_S\} \quad \tilde{\sigma}(\boldsymbol{\pi_T}) = \{s \mid \forall \mathbf{t}.\, s = \mathbf{t}|_{\Sigma_S} \Rightarrow \mathbf{t} \in \boldsymbol{\pi_T}\}$$

That is, the target guarantee of a source property is that the target has the same source-level behavior, sprinkled with arbitrary target-level behavior. Conversely, the source-level obligation of a target property is the aggregate of those source traces all of whose target-level enrichments are in the target property.

Since R^S and $\mathbf{R^T}$ are very similar, it is simple to prove that the identity compiler $(\cdot\downarrow)$ from R^S to $\mathbf{R^T}$ is secure according to the trace relation $\sim$ defined above.

Theorem 5.4 **($\cdot\downarrow$ is secure).** $\cdot\downarrow$ *is* $\mathsf{RTC}^{\sim}$

5.3 Instances of Trace-Relating Robust Preservation of Safety and Hypersafety

To provide examples of cross-language trace-relations that preserve safety and hypersafety properties, we show how existing secure compilation results can be interpreted in our framework. This indicates how the more general theory developed here can already be instantiated to encompass existing results, and that existing proof techniques can be used in order to achieve the secure compilation criteria we define.

Concerning the preservation of safety, Patrignani and Garg [37] consider a compiler from a typed, concurrent WHILE language to an untyped, concurrent WHILE language with support for memory capabilities. As in §3.3, their source has bools and nats while their target only has **nat**s. Additionally, their source has an ML-like memory model (where the domain of the heap is arbitrary locations ℓ) while their target has an assembly-like memory model (where the domain of the heap is natural numbers $\mathbf{n}$). Their traces consider context-program interactions and as such they are concatenations of call and return actions with parameters, which can include booleans as well as locations. Because of the aforementioned differences, they need a cross-language relation to relate source and target actions.

Besides defining a relation on traces (i.e., an instance of $\sim$), they also define a relation between source and target safety properties. They provide an instantiation of τ that maps all safe source traces to the related target ones. This ensures that no additional target trace is introduced in the target property, and source safety properties are mapped to target safety ones by τ. Their compiler is then proven to generate code that respects τ, so they achieve a variation of $\mathsf{RTP}^{Safe \circ \tilde{\tau}}$.

Concerning the preservation of hypersafety, Patrignani and Garg [36] consider compilers in a reactive setting where traces are concatenations of input ($\alpha?$) and output ($\alpha!$) actions. In their setting, traces are different between source and target, so they define a cross-language relation on actions that is total on the source actions and injective.

Additionally, their set of target output actions is strictly larger than the source one, as it includes a special action $\surd$, which is how compiled code must respond to invalid target inputs (i.e., receiving a $\mathbf{bool}$ when a $\mathbf{nat}$ was expected). Starting from the relation on actions, they define $\mathbf{TPC}$, which is an instance of what we call τ. Informally, given a set of source traces, $\mathbf{TPC}$ generates all target traces that are related (pointwise) to a source trace. Additionally, it generates all traces with interleavings of undesired inputs $\alpha?$ followed by $\surd$ as long as removing $\alpha?\surd$ leaves a trace that relates the source trace.

$\mathbf{TPC}$ preserves hypersafety across languages, i.e., it is an instance of $\mathsf{RSCHP}^{HSafe \circ \tilde{\tau}}$ mapping source hypersafety to target hypersafety (and safety to safety).

6 Related Work

We already discussed how our results relate to some existing work in correct compilation [27, 45] and secure compilation [2, 36, 37]. We explicitly notice that most of our definitions and results make no assumptions about the structure of traces. Referring to traces as sequences of events is only needed to define safety properties and hypersafety. In all other cases, traces are arbitrary and could for instance be "interaction trees" [49]. In fact, we believe that the compilation from IMP to ASM proposed by Xia et al. [49] can be seen as an instace of $\mathsf{HC}^{\sim}$, for the relation they call "trace equivalence."

Compilers where our work could be useful. Our work should be broadly applicable to understanding the guarantees provided by many verified compilers. For instance, Wang et al. [48] recently proposed a CompCert variant that compiles all the way down to machine code, and it would be interesting to see if the model at the end of §3.1 applies there too. This and, many other verified compilers [9, 23, 33, 43] beyond CakeML [45] deal with resource exhaustion and it would be interesting to also apply the ideas of §3.2 to them. Hur and Dreyer [21] devised a correct compiler from an ML language to assembly using a cross-language logical relation to state their CC theorem. They do not have traces, though were one to add them, the logical relation on values would serve as the basis for the trace relation and therefore their result would attain $\mathsf{CC}^{\sim}$.

Switching to more informative traces capturing the interaction between the program and the context is often used as a proof technique for secure compilation [2, 22, 35]. Most of these results consider a cross-language relation, so they probably could be proved to attain one of the criteria from Figure 2.

Generalizations of compiler correctness. The compiler correctness definition of Morris [32] was already general enough to account for trace relations, since it considered a translation between the semantics of the source program and that of the compiled program, which he called "decode" in his diagram, reproduced in Figure 3 (left). And even some of the more recent compiler correctness definitions preserve this kind of flexibility [38]. Still, while $\mathsf{CC}^{\sim}$ can be seen as an instance of Morris' [32] definition, we are not aware of any prior work that investigated the preservation of properties, when the "decode translation" is neither the identity nor a bijection, and source properties need to be re-interpreted as target ones and vice versa.

Correct compilation and Galois connections. Melton et al. [30] and Sabry and Wadler [42] expressed a strong variant of compiler correctness using the diagram of Figure 3

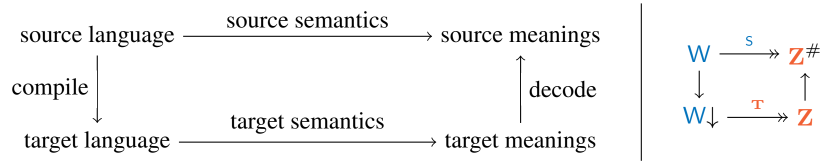

Fig. 3: Morris's [32] (left) and Melton et al.'s [30] and Sabry and Wadler's [42] (right) correctness diagrams.

(right) [30, 42]. They require that compiled programs *parallel* the computation steps of the original source programs, fact that can be proven showing the existence of a *decompilation* map $\#$ that makes the diagram commute, or equivalently, the existence of an adjoint for $\downarrow$ ($W \leq W' \iff W \twoheadrightarrow W'$ for both source and target). The "parallel" intuition can be formalized as an instance of $\mathsf{CC}^{\sim}$. Take source and target traces to be finite or infinite sequences of program states (maximal trace semantics [12]), and relate them exactly like Melton et al. [30] and Sabry and Wadler [42].

7 Conclusion and Future Work

We have extended the property preservation view on compiler correctness to arbitrary trace relations, and believe that this will be useful for understanding the guarantees various compilers provide. An open question is whether, given a compiler, there exists a most precise $\sim$ relation for which this compiler is correct. As mentioned in §1, every compiler is $\mathsf{CC}^{\sim}$ for some $\sim$, but under which conditions there is a most precise relation? In practice, more precision may not always be better though, as it may be at odds with compiler efficiency and may not align with more subjective notions of the usefulness, leading to tradeoffs in the selection of suitable relations. Finally, another interesting direction for future work is studying whether using the relation to Galois connections allows us to more easily compose trace relations for different purposes, say, for a compiler whose target language has undefined behavior, resource exhaustion, and side-channels. In particular, are there ways to obtain complex relations by combining simpler ones in a way that eases the compiler verification burden?

Bibliography

[1] M. Abadi, A. Banerjee, N. Heintze, N. Heintze, and J. G. Riecke. A core calculus of dependency. In *Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. 1999.

[2] C. Abate, R. Blanco, D. Garg, C. Hriţcu, M. Patrignani, and J. Thibault. Journey beyond full abstraction: Exploring robust property preservation for secure compilation. *CSF*, 2019.

[3] A. Ahmed, D. Garg, C. Hriţcu, and F. Piessens. Secure Compilation (Dagstuhl Seminar 18201). *Dagstuhl Reports*, 8(5), 2018.

[4] J. B. Almeida, M. Barbosa, G. Barthe, B. Grégoire, A. Koutsos, V. Laporte, T. Oliveira, and P. Strub. The last mile: High-assurance and high-speed cryptographic implementations. *CoRR*, abs/1904.04606, 2019.

[5] A. Anand, A. Appel, G. Morrisett, Z. Paraskevopoulou, R. Pollack, O. S. Belanger, M. Sozeau, and M. Weaver. CertiCoq: A verified compiler for Coq. CoqPL Workshop, 2017.

[6] G. Barthe, B. Grégoire, and V. Laporte. Secure compilation of side-channel countermeasures: the case of cryptographic "constant-time". *CSF*. 2018.

[7] L. Beringer, G. Stewart, R. Dockins, and A. W. Appel. Verified compilation for shared-memory C. In Z. Shao, editor, *Programming Languages and Systems - 23rd European Symposium on Programming, ESOP 2014, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2014, Grenoble, France, April 5-13, 2014, Proceedings*. 2014.

[8] F. Besson, S. Blazy, and P. Wilke. A verified compcert front-end for a memory model supporting pointer arithmetic and uninitialised data. *J. Autom. Reasoning*, 62(4), 2019.

[9] Q. Carbonneaux, J. Hoffmann, T. Ramananandro, and Z. Shao. End-to-end verification of stack-space bounds for C programs. In M. F. P. O'Boyle and K. Pingali, editors, *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '14, Edinburgh, United Kingdom - June 09 - 11, 2014*. 2014.

[10] C. Cimpanu. Microsoft: 70 percent of all security bugs are memory safety issues. ZDNet, 2019.

[11] M. R. Clarkson and F. B. Schneider. Hyperproperties. *JCS*, 18(6), 2010.

[12] P. Cousot. Constructive design of a hierarchy of semantics of a transition system by abstract interpretation. *Theoretical Computer Science*, 277(1-2), 2002.

[13] P. Cousot and R. Cousot. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. *POPL*. 1977.

[14] V. D'Silva, M. Payer, and D. X. Song. The correctness-security gap in compiler optimization. *S&P Workshops*. 2015.

[15] R. Giacobazzi and I. Mastroeni. Abstract non-interference: a unifying framework for weakening information-flow. *ACM Transactions on Privacy and Security (TOPS)*, 21(2), 2018.

[16] J. A. Goguen and J. Meseguer. Security policies and security models. *S&P*, 1982.

[17] R. Gu, Z. Shao, J. Kim, X. N. Wu, J. Koenig, V. Sjöberg, H. Chen, D. Costanzo, and T. Ramananandro. Certified concurrent abstraction layers. In J. S. Foster and D. Grossman, editors, *Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2018, Philadelphia, PA, USA, June 18-22, 2018*. 2018.

[18] I. Haller, Y. Jeon, H. Peng, M. Payer, C. Giuffrida, H. Bos, and E. van der Kouwe. TypeSan: Practical type confusion detection. *CCS*, 2016.

[19] Heartbleed. The Heartbleed bug. `http://heartbleed.com/`, 2014.

[20] C. Hriţcu, , D. Chisnall, D. Garg, and M. Payer. Secure compilation. SIGPLAN PL Perspectives Blog, 2019.

[21] C.-K. Hur and D. Dreyer. A Kripke logical relation between ML and Assembly. *SIGPLAN Not.*, 46(1), 2011.

[22] A. Jeffrey and J. Rathke. Java Jr: Fully abstract trace semantics for a core Java language. *ESOP*. 2005.

[23] J. Kang, C. Hur, W. Mansky, D. Garbuzov, S. Zdancewic, and V. Vafeiadis. A formal C memory model supporting integer-pointer casts. *PLDI*, 2015.

[24] J. Kang, Y. Kim, C.-K. Hur, D. Dreyer, and V. Vafeiadis. Lightweight verification of separate compilation. *POPL*, 2016.

[25] L. Lamport and F. B. Schneider. Formal foundation for specification and verification. In *Distributed Systems: Methods and Tools for Specification, An Advanced Course*, 1984.

[26] C. Lattner. What every C programmer should know about undefined behavior #1/3. LLVM Project Blog, 2011.

[27] X. Leroy. Formal verification of a realistic compiler. *CACM*, 52(7), 2009.

[28] X. Leroy. The formal verification of compilers (DeepSpec Summer School 2017), 2017.

[29] J. McCarthy and J. Painter. Correctness of a compiler for arithmetic expressions. *Mathematical Aspects Of Computer Science 1*, 19 of Proceedings of Symposia in Applied Mathematics, 1967.

[30] A. Melton, D. A. Schmidt, and G. E. Strecker. Galois connections and computer science applications. In *Proceedings of a Tutorial and Workshop on Category Theory and Computer Programming*. 1986.

[31] R. Milner and R. Weyhrauch. Proving compiler correctness in a mechanized logic. In *Proceedings of 7th Annual Machine Intelligence Workshop, volume 7 of Machine Intelligence*. 1972.

[32] F. L. Morris. Advice on structuring compilers and proving them correct. *POPL*. 1973.

[33] E. Mullen, D. Zuniga, Z. Tatlock, and D. Grossman. Verified peephole optimizations for CompCert. *PLDI*, 2016.

[34] G. Neis, C. Hur, J. Kaiser, C. McLaughlin, D. Dreyer, and V. Vafeiadis. Pilsner: a compositionally verified compiler for a higher-order imperative language. *ICFP*. 2015.

[35] M. Patrignani and D. Clarke. Fully abstract trace semantics for protected module architectures. *CL*, 42, 2015.

[36] M. Patrignani and D. Garg. Secure compilation and hyperproperty preservation. *CSF*. 2017.

[37] M. Patrignani and D. Garg. Robustly safe compilation. *ESOP*, 2019.

[38] D. Patterson and A. Ahmed. The next 700 compiler correctness theorems (functional pearl). To appear at ICFP, 2019.

[39] T. Ramananandro, Z. Shao, S. Weng, J. Koenig, and Y. Fu. A compositional semantics for verified separate compilation and linking. *CPP*. 2015.

[40] J. Regehr. A guide to undefined behavior in C and C++, part 3. Embedded in Academia blog, 2010.

[41] A. Sabelfeld and D. Sands. Dimensions and principles of declassification. In *Computer Security Foundations 18th Workshop*. 2005.

[42] A. Sabry and P. Wadler. A reflection on call-by-value. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 19(6), 1997.

[43] J. Sevcík, V. Vafeiadis, F. Z. Nardelli, S. Jagannathan, and P. Sewell. Relaxed-memory concurrency and verified compilation. *POPL*. 2011.

[44] G. Stewart, L. Beringer, S. Cuellar, and A. W. Appel. Compositional CompCert. *POPL*. 2015.

[45] Y. K. Tan, M. O. Myreen, R. Kumar, A. Fox, S. Owens, and M. Norrish. The verified CakeML compiler backend. *Journal of Functional Programming*, 29, 2019.

[46] X. Wang, H. Chen, A. Cheung, Z. Jia, N. Zeldovich, and M. F. Kaashoek. Undefined behavior: What happened to my code? *APSYS*, 2012.

[47] X. Wang, N. Zeldovich, M. F. Kaashoek, and A. Solar-Lezama. Towards optimization-safe systems: Analyzing the impact of undefined behavior. *SOSP*. 2013.

[48] Y. Wang, P. Wilke, and Z. Shao. An abstract stack based approach to verified compositional compilation to machine code. *PACMPL*, 3(POPL), 2019.

[49] L.-y. Xia, Y. Zakowski, P. He, C.-K. Hur, G. Malecha, B. C. Pierce, and S. Zdancewic. Interaction trees: Representing recursive and impure programs in coq (work in progress). *arXiv preprint arXiv:1906.00046*, 2019.

[50] J. Zhao, S. Nagarakatte, M. M. K. Martin, and S. Zdancewic. Formalizing the LLVM intermediate representation for verified program transformations. *POPL*. 2012.